

# From Numerical Cosmology to Efficient Bit Abstractions for the Standard Library


Vincent Reverdy

Department of Astronomy  
University of Illinois at Urbana-Champaign (UIUC)

September 21, 2016

# Prologue

# What is this talk about?



*Levels of parallelism from Bryce's talk*

## Alternative titles

- Bit-level parallelism (manipulating bit sequences by chunks)
- Abstracting bits
- On bit manipulation algorithms for the standard library
- What's wrong with you `std::vector<bool>`?
- Reinventing `std::bitset` and `std::vector<bool>`
- Counting bits 100× faster than with `std::vector<bool>`
- Playing with 0 and 1

# Using bit utilities in less than 2 minutes (1/2)

## Step 1: Clone and download

- With GIT: `git clone https://github.com/vreverd/bit.git`
- Without GIT: go on the page <https://github.com/vreverd/bit>, click on Clone or download → Download ZIP, download, and unzip it.

## Step 2: Run a minimal test case

```
1 #include <iostream>
2 #include "../bit/cpp/bit.hpp" // Your path to bit.hpp
3 using namespace bit;
4
5 int main(int argc, char* argv[]) {
6     using uint_t = unsigned int;
7     uint_t n = 42;
8     auto first = bit_iterator<uint_t*>(&n);
9     auto last = bit_iterator<uint_t*>(&n + 1);
10    for (; first != last; ++first) std::cout<<*first;
11    std::cout<<std::endl;
12    return 0;
13 }
14
15 // Compilation with GCC: g++ -std=c++14 -pedantic -O3 main.cpp -o main
16 // Output: 01010100000000000000000000000000
```

Compile it and run it, it should display the bits of `n` from the LSB to the MSB.



# Using bit utilities in less than 2 minutes (2/2)

## Step 3: And that's it

You are all set...

## Contact and links

- GITHUB: <https://github.com/vreverdy/bit>
- Contact: [vince.rev@gmail.com](mailto:vince.rev@gmail.com)
- ISO C++ proposal (description): P0237R0
- ISO C++ proposal (last wording): P0237R2 (in progress)
- Collaborations are very welcome! Same for comments! Same for benchmarks!

# Chapter I: An astrophysical motivation

Once upon a time...



...in a galaxy far far away...

...on a small piece of rock...



...wandering aimlessly in a vast Universe...



...a team of astrophysicists was wondering  
about the nature of life, the Universe, and everything.

Because they knew some maths, some physics  
some computer science, and some programming,


```
typename std::enable_if<std::is_integer<T>::value>::type
```

$$G_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}$$

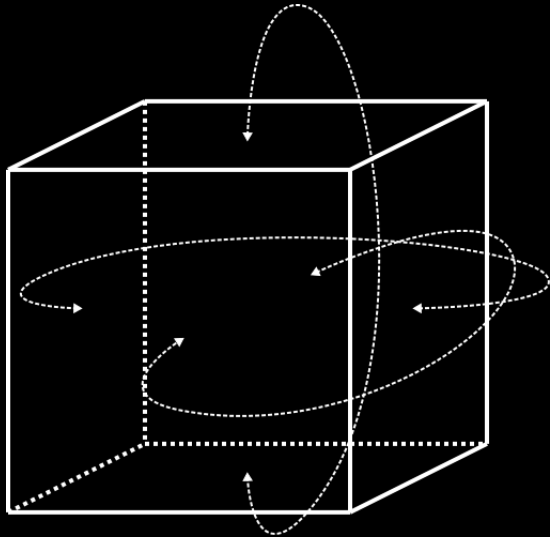
$$\nabla_r^2 \phi = 4\pi G\rho$$

$$\dot{\rho} = -3H\left(\rho + \frac{p}{c^2}\right)$$

they decided to design a code that could answer  
their (meta)physical questions.




They said: "Let's take an enormous box..."



...with periodic boundary conditions...  
(right/left, top/bottom, and front/back connected together)









...and let's fill that enormous box with particles weighing the mass of millions or billions of suns...

(note: yes, that's kind of huge)



Now, divide the box in cells using a regular grid and apply the following recipe:



1) For each cell  $c$  containing particles with position  $\vec{x}$  and velocity  $\vec{v}$

2) Interpolate the density  $\rho$  in the cell  $c$  depending on surrounding particles

3) From  $\rho$ , compute the gravitational potential  $\Phi$


4) From  $\Phi$ , interpolate back the acceleration  $\vec{a}$  at the position of particles

5) From  $\vec{a}$ , compute the new speed  $\vec{v}$  of each particle

6) From  $\vec{v}$ , compute the new position  $\vec{x}$  of each particle

7) Restart at 1) with the updated position  $\vec{x}$  and speed  $\vec{v}$





Using this recipe with millions of particles  
we can simulate galaxy formation!"

Simulating galaxies is nice, but it was not answering their (meta)physical questions.

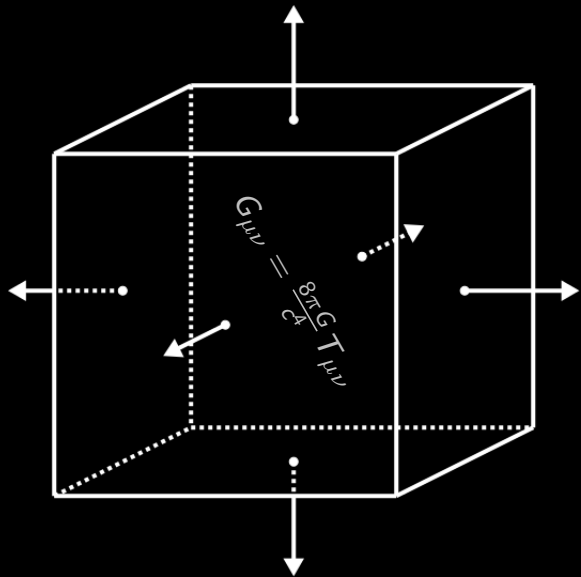


So they decided to do better. "Let's try to investigate larger scales with galaxy-sized particles" they said.



First, they took a supercomputer.







Second, they made the box expand as the Universe does\* according to the theory of General Relativity (GR)

(\*here, they did not consider backreaction effects)






measured cosmological  
microwave background




Third, they filled the box with billions of particles with the same distribution (statistically speaking) as the matter in the primordial Universe



Quadtree (2D)



Octree (3D)




Fourth, they updated their algorithm using an Adaptive Mesh Refinement (AMR) strategy to increase the resolution in regions of interest.

And finally, after all this work, they ran their simulation, using millions of computing hours over thousands of cores and generating hundreds of terabytes of data.


And this is what they obtained.



DEUS  
DEUS  
DEUS  
DEUS



Actual cosmic  
web structure  
observed by the SDSS





Simulated cosmic web structure  
(each point is of the mass of a Milky Way)




It was nice and exciting,

But to answer their questions,  
they needed far more computing power.

Thankfully, new architectures and supercomputers were coming...



But they soon realized that there was a MAJOR problem:  
their code would not scale up to millions of cores!



And a part of this problem can be boiled down to data structures:  
 on supercomputers, pure computing capabilities are improving faster  
 than memory performances

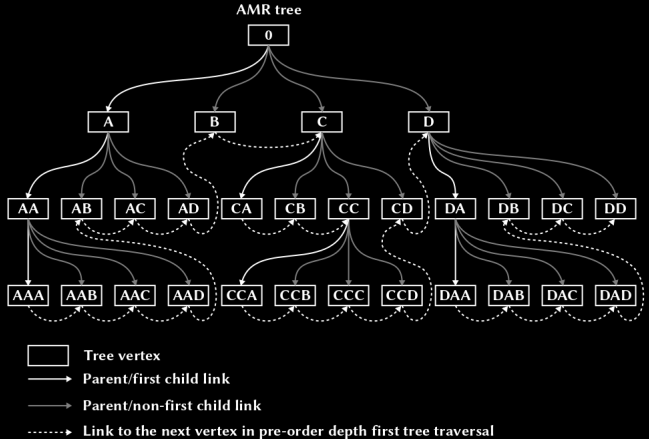

|                                    |             |    |      |                        |
|------------------------------------|-------------|----|------|------------------------|
| One cycle on a 3 GHz processor     | 1           | ns |      |                        |
| L1 cache reference                 | 0.5         | ns |      |                        |
| Branch mispredict                  | 5           | ns |      |                        |
| L2 cache reference                 | 7           | ns |      | 14x L1 cache           |
| Mutex lock/unlock                  | 25          | ns |      |                        |
| Main memory reference              | 100         | ns |      | 20x L2, 200x L1        |
| Compress 1K bytes with Snappy      | 3,000       | ns |      |                        |
| Send 1K bytes over 1 Gbps network  | 10,000      | ns | 0.01 | ms                     |
| Read 4K randomly from SSD*         | 150,000     | ns | 0.15 | ms                     |
| Read 1 MB sequentially from memory | 250,000     | ns | 0.25 | ms                     |
| Round trip within same datacenter  | 500,000     | ns | 0.5  | ms                     |
| Read 1 MB sequentially from SSD*   | 1,000,000   | ns | 1    | ms 4X memory           |
| Disk seek                          | 10,000,000  | ns | 10   | ms 20x datacenter RT   |
| Read 1 MB sequentially from disk   | 20,000,000  | ns | 20   | ms 80x memory, 20X SSD |
| Send packet CA->Netherlands->CA    | 150,000,000 | ns | 150  | ms                     |

From Chandler Carruth "Efficiency with Algorithms, Performance with Data Structures" (cppcon 2014)

Credit: Jeffrey Dean, Google Research



Explicit trees kill performances because of poor cache-awareness



So they decided to get rid of explicit trees to make the most of supercomputers.

Numerical  
cosmology








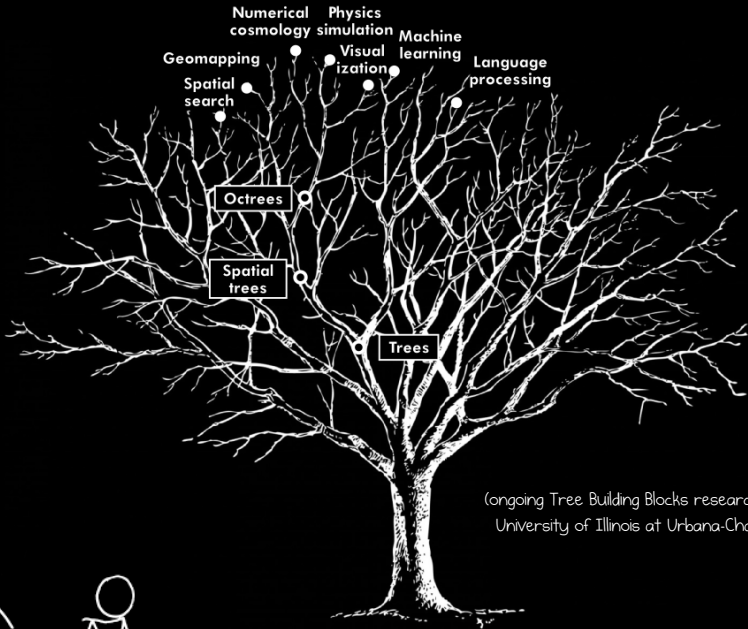
Numerical Physics  
cosmology simulation

Visualization

Octrees


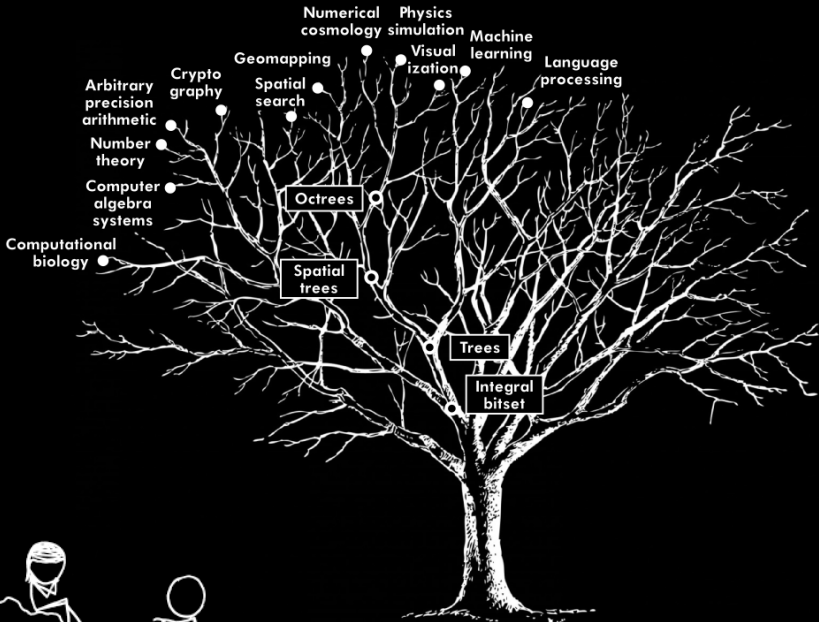


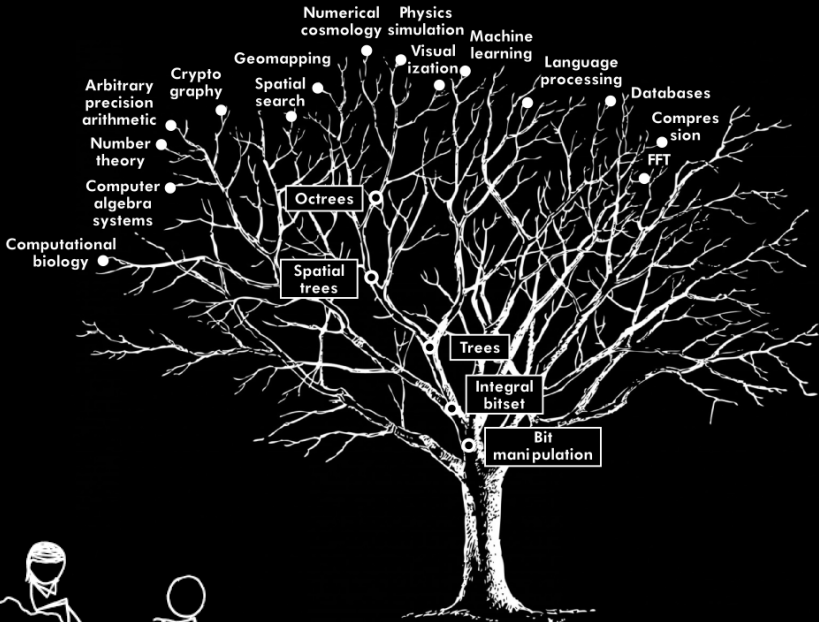





(ongoing Tree Building Blocks research project,  
University of Illinois at Urbana-Champaign)









They soon realized that to make the most efficient trees ever they would need very efficient ways to manipulate bits.

And this is how this whole story started...



They soon realized that to make the most efficient trees ever they would need very efficient ways to manipulate bits.

And this is how this whole story started...

Thankfully, the old magician of C++ lands came to them to help them to start their quest





They soon realized that to make the most efficient trees ever they would need very efficient ways to manipulate bits.

And this is how this whole story started...

Thankfully, the old magician of C++ lands came to them to help them to start their quest

He gave them a very precious `(std::)map` and disappeared...






TERRA INCOGNITA

STL

BOOST

external  
libraries  
seas



Bitwise canyon

THE HOLY STANDARD TERRITORY

Memory management desert

The C++ Lands

Its amazing creatures and weird beasts

explorer  
Aleria (<http://alernacpp.blogspot.com/>)

cartographer  
Jim (<http://jimblog.me/>)

2012 edition




Loki

OL

ACC



TERRA INCOGNITA



THE HOLY STANDARD TERRITORY



# The C++ Lands

Its amazing creatures and weird beasts


explorer  
Aleria (<http://aleniaccpp.blogspot.com/>)  
2012 edition

cartographer  
Jim (<http://jimblog.me/>)



ACC

So they decided to go there  
with the hope to leverage bit manipulation  
to speed up their trees and their simulations...



## Chapter II: Manipulating bits with the current ISO C++ Standard

# Constant-size bit container: `std::bitset`

```
template <std::size_t N> class bitset
```

The class template `bitset` represents a fixed-size sequence of `N` bits. Bitsets can be manipulated by standard logic operators and converted to and from strings and integers. (*source: [cppreference.com](http://cppreference.com)*)

```
1 // Example from cppreference.com
2 #include <iostream>
3 #include <bitset>
4
5 int main() {
6     // Initialization
7     std::bitset<8> b("00010010");
8     std::cout<<"initial value: "<<b<<'\n';
9     // Find the first unset bit
10    size_t idx = 0;
11    while (idx < b.size() && b[idx]) ++idx;
12    // Continue setting bits until half the bitset is filled
13    while (idx < b.size() && b.count() < b.size()/2) {
14        b.set(idx);
15        std::cout<<"setting bit "<<idx<<": "<<b<<'\n';
16        while (idx < b.size() && b.test(idx)) ++idx;
17    }
18 }
19
20 // Output
21 // initial value: 00010010
22 // setting bit 0: 00010011
23 // setting bit 2: 00010111
```

# Constant size bit container: `std::bitset`

## What is good with bitsets?

- Very simple and easy to use
- Set of optimized members like `test`, `all`, `any`, `none`, `count`, `set`, `reset`, `flip`

## Limitations

- Very limited functionality
- No `begin` and `end` iterators: not compatible with algorithms and standard containers
- No control on the underlying representation

## Underlying representation

In terms of implementation, bits are likely to be stored in a contiguous array of unsigned integers. However, there is no way to access this underlying representation.



# What does `std::bitset::operator[]` return?

```
1 std::bitset<8> b(42);
2 bool boolean = b[0];
3 auto something = b[0];
4 boolean += 1;
5 something += 1; // Failure
```

## Limitations

- `std::bitset::operator[]` returns a proxy of type `std::bitset::reference`
- *Almost* like a `bool`...
- But only *almost*: different promotion rules, different arithmetic, a member flip and a different behavior for `operator~`
- Very confusing and error-prone

## Very confusing and error-prone?

```
1 // Initialization
2 std::bitset<8> b(0);
3 bool boolean = b[0];
4 auto something = b[0];
5
6 // Operations
7 boolean = ~(~boolean);
8 something = ~(~something);
9
10 // Display
11 std::cout<<boolean<<something<<std::endl;
```

### Question

What does that print? 00, 11, 10 or 01?

# Very confusing and error-prone?

```
1 // Initialization
2 std::bitset<8> b(0);
3 bool boolean = b[0];
4 auto something = b[0];
5
6 // Operations
7 boolean = ~(~boolean);
8 something = ~(~something);
9
10 // Display
11 std::cout<<boolean<<something<<std::endl;
```

## Question

What does that print? 00, 11, 10 or 01?

## Answer

10

## Why?

- `boolean` → `~false` → `~(-1)` → `boolean = 0` → `boolean == 0`
- `something` → `~reference(0)` → `~(true)` → `something = -2` → `something == 1`

# Dynamic size bit container: `std::vector<bool>`

## What is `std::vector<bool>`?

A mistake


```
template <class Allocator> class vector<bool, Allocator>
```

`std::vector<bool>` is a space-efficient specialization of `std::vector` for the type `bool`. The manner in which `std::vector<bool>` is made space efficient (as well as whether it is optimized at all) is implementation defined. One potential optimization involves coalescing vector elements such that each element occupies a single bit instead of `sizeof(bool)` bytes. `std::vector<bool>` behaves similarly to `std::vector`, but in order to be space efficient, it:

- Does not necessarily store its elements as a contiguous array (so `&v[0] + n != &v[n]`)
- Exposes class `std::vector<bool>::reference` as a method of accessing individual bits. In particular, objects of this class are returned by `operator[]` by value.
- Does not use `std::allocator_traits::construct` to construct bit values.
- Does not guarantee that different elements in the same container can be modified concurrently by different threads.

(source: [cppreference.com](http://cppreference.com))

# Dynamic size bit container: `std::vector<bool>`



## What is `std::vector<bool>`?

It's not a container

## What's wrong with you `std::vector<bool>`?

- “`std::vector<bool>` is nonconforming, and forces optimization choice”, H. Sutter (N1185) [1999]
- “`std::vector<bool>` : More problems, better solutions”, H. Sutter (N1847) [2005]
- “Library issue 96: Fixing `std::vector<bool>`”, B. Dawes (N2160) [2007]
- “A specification to deprecate `std::vector<bool>`”, A. Meredith (N2204) [2007]

⇒ 2016: `std::vector<bool>` is still alive

# Dynamic size bit container: `std::vector<bool>`

## What is good with vector bool?

- Compact storage in memory
- begin and end iterators: compatible with standard algorithms

## Problems

- Poor performances
- No access to the underlying representation
- No thread safety
- Breaks the normal behavior of containers
- Error-prone behavior of `std::vector<bool>::reference` (*almost a bool*).

## On dynamic bitsets

The functionality is ok, but specializing `std::vector` for it was not the best idea ever. A `std::dynamic_bitset` (as in BOOST) would have been a far better option. But even with that, most of the problems remain.

# Bit manipulation

## Example of use of bit instructions (Hilewitz, 2008)

| Application \ Instruction           |                     | bfly | pex | pdep | setib | setb | pex.v | pdep.v | grp | bmm |
|-------------------------------------|---------------------|------|-----|------|-------|------|-------|--------|-----|-----|
| Binary Compression                  |                     |      | ×   |      |       |      |       |        |     | (×) |
| Binary Decompression                |                     |      |     | ×    |       |      |       |        |     | (×) |
| LSB                                 | Encoding            |      |     | ×    |       | ×    |       |        |     | (×) |
|                                     | Decoding            |      | ×   |      | ×     |      |       |        |     | (×) |
| Steganography                       | Encoding            |      | ×   |      | ×     |      |       |        |     | (×) |
|                                     | Decoding            |      |     | ×    |       |      |       |        |     | (×) |
| Transfer Coding                     | Encoding            |      | ×   |      |       |      |       |        |     | (×) |
|                                     | Decoding            |      |     | ×    |       |      |       |        |     | (×) |
| Integer                             | Compression         |      |     | ×    |       |      |       |        |     | (×) |
|                                     | Decompression       |      | ×   |      |       |      |       |        |     | (×) |
| Binary Image Morphology             |                     |      | ×   |      |       |      | (×)   |        |     | (×) |
| Random Number Generation            | Von Neumann         |      |     |      |       |      | ×     |        |     |     |
|                                     | Toeplitz            |      |     |      |       |      |       |        |     | ×   |
| Bioinformatics                      | Compression         |      | ×   |      |       |      |       |        |     | (×) |
|                                     | BLASTZ Alignment    |      | ×   |      |       |      |       |        |     | (×) |
|                                     | BLASTX Translation  |      |     | ×    |       |      |       |        |     | (×) |
|                                     | Reversal            | ×    |     |      |       |      |       |        |     | (×) |
| Cryptography                        | Block Ciphers       | ×    |     |      |       |      |       |        | ×   | ×   |
|                                     | Stream Ciphers      |      |     |      |       |      | ×     |        |     | ×   |
|                                     | Public Key          |      |     |      |       |      |       |        |     | ×   |
|                                     | Future              | ?    | ?   | ?    | ?     | ?    | ?     | ?      | ?   | ?   |
| Cryptanalysis                       | Linear              |      |     |      |       |      |       |        |     | ×*  |
|                                     | Algebraic           |      |     |      |       |      |       |        |     | ×   |
|                                     | Future/Proprietary  | ?    | ?   | ?    | ?     | ?    | ?     | ?      | ?   | ?   |
| DARPA HPCS Discrete Math Benchmarks | Matrix Transpose    |      |     |      |       |      |       |        |     | ×*  |
|                                     | Equation Solving    |      |     |      |       |      |       |        |     | ×   |
| Linear Feedback Shift Registers     |                     |      |     |      |       |      |       |        |     | ×   |
| Error Correction                    | Block Codes         |      |     |      |       |      |       |        |     | ×   |
|                                     | Convolutional Codes |      |     |      |       |      |       |        |     | ×   |
|                                     | Puncturing          |      | ×   | ×    |       |      |       |        |     | (×) |

# The current state of bit manipulation

## Instructions

CPUs include more and more very efficient bit manipulation instructions but they are often left unused because of the lack of standard utilities to access them.

## Bit Manipulation Instruction Sets

- ABM: Advanced Bit Manipulation, POPCNT, LZCNT (Intel SSE 4.2, AMD ABM)
- BMI1: Bit Manipulation Instruction Set 1 ( $\geq$  Intel Haswell, AMD Piledriver)
- BMI2: Bit Manipulation Instruction Set 2 ( $\geq$  Intel Haswell, AMD Excavator)
- TBM: Trailing Bit Manipulation ( $\geq$  AMD Piledriver)
- BME: Bit Manipulation Engine (ARM Cortex)

## Compiler intrinsic examples

- `_bzhi_u32, _bzhi_u64, _pdep_u32, _pdep_u64, _pext_u32, _pext_u64`
- `__builtin_clz, __builtin_ctz, __builtin_clrsb, __builtin_popcount`
- `__builtin_ia32_bextr_u32, __builtin_ia32_bextr_u64`
- `__builtin_ia32_lzcnt_u32, __builtin_ia32_lzcnt_u64`



# The current state of bit manipulation

## Past proposal

“A constexpr bitwise operations library for C++” (Fioravante, 2014)

## Problems

- More than 60 bit-specialized functions: too many functions
- Too low level and domain specific
- Limited to integral arguments
- Does not provide tools to manipulate arbitrary long sequence of bits

## Genericity

Need of something far more generic

## Chapter III: Introducing The Bit Library

# Summary of The Bit Library and P0237

## What?

Designing tools to provide a generic way to operate on bits and sequence of bits.

## Why?

- To be able to use unsigned integers as sets of bits
- To make the most of bit manipulation instruction sets
- To provide users with utilities to build fast bit manipulation algorithms
- To provide users with utilities to build efficient bit-based data structures

## Application areas

- |                    |                            |                                  |
|--------------------|----------------------------|----------------------------------|
| ■ Hashing          | ■ Random number generation | ■ High-performance computing     |
| ■ Video games      | ■ Binary compression       | ■ Arbitrary-precision arithmetic |
| ■ Image processing | ■ Error correction         |                                  |
| ■ Cryptography     |                            |                                  |

## Concrete examples of use

- User-defined bit sets and bit arrays, iteration over bits
- Access to the underlying bits of bounded and unbounded integers
- Allowing `std::count` to call a POPCNT instruction when executed on bits

# Motivation

## Bjarne Stroustrup - The C++ Programming Language (2013)

*“unsigned integer types are ideal for uses that treat storage as a bit array.”*

### Except that...

... there is no standard way to access and manipulate bits in C++.

### Main motivations

- Ease the use of unsigned integers as bit arrays
- Provide a standard way to access and manipulate unique bits (set, reset, flip...)
- Provide an abstraction to leverage bit manipulation instruction sets
- Provide efficient versions of standard algorithms on bits
- Facilitate the design of fast bit manipulation algorithms on sequences of bits
- Facilitate the implementation of data structures based on bit sequences
- Provide tools to access the underlying representation of such data structures
- Provide tools to easily build alternatives to `std::vector<bool>`

### Objectives

Simplicity, genericity and efficiency.

## Solution summary

### Key idea: `bit_iterator`

All of that can be achieved through a carefully designed `bit_iterator` acting as an iterator adaptor (like `std::reverse_iterator`).

```
template <class Iterator> class bit_iterator
```

A `bit_iterator` is a tool that allows to reinterpret a sequence of unsigned integers as a sequence of bits. It provides an API that is both:

- high-level: easy to use from the general user point of view
- low-level: gives access to the underlying representation for optimization purposes

### High-level point of view: counting bits (assuming `std::` prefix)

```
1 // Initialization
2 using uint_t = unsigned int;
3 using container_t = std::list<uint_t>;
4 using iterator_t = typename container_t::iterator;
5 container_t container = {0, 1, 2, 3, 4};
6
7 // From the bit number 5 of container[0], to the end of the container
8 std::bit_iterator<iterator_t> first(std::begin(container), 5);
9 std::bit_iterator<iterator_t> last(std::end(container));
10
11 // Counts the bits set to 1
12 std::cout<<std::count(first, last, std::one_bit)<<std::endl;
```

## Solution summary

```
unsigned int i = 1751901291;
```

Object representation in memory (little endian)

```
01101011 11100100 01101011 01101000
01101000 00011010 10111110 01000110 1011
```

Binary value

```
std::list<unsigned int> uintlist =
{1751901291, 1751901292, 1751901293, 1751901294};
```

```
31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16 15 14 13 12 11 10 09 08 07 06 05 04 03 02 01 00
00 0110101000 01101011111110010001101011
01 0110100000110101111110010001101100
02 0110100000110101111110010001101101
03 0110100000110101111110010001101110
```

■ - 23 = ■

■ + 55 = ■

# Solution summary

## 3 key functions

- `it.base()` (`bit_iterator::base`): returns the underlying iterator, inspired from `std::reverse_iterator::base`
- `it.position()` and `(*it).position()` (`bit_iterator::position` and `bit_reference::position`): returns the current bit position in the underlying value, starting from 0 (LSB) to `binary_digits<T>::value - 1` (MSB)
- `(*it).address()` (`bit_reference::address`): returns a pointer to the underlying value

## Low-level point of view: counting bits (1/2)

```
1  template <class InputIt>
2  typename bit_iterator<InputIt>::difference_type
3  count(
4      bit_iterator<InputIt> first,
5      bit_iterator<InputIt> last,
6      bit_value value
7  )
8  {
9      // Assertions
10     _assert_range_viability(first, last);
11
12     // Initialization
13     using underlying_type = typename bit_iterator<InputIt>::underlying_type;
14     using difference_type = typename bit_iterator<InputIt>::difference_type;
15     constexpr difference_type digits = binary_digits<underlying_type>::value;
```

# Solution summary

## Low-level point of view: counting bits (2/2)

```
1  difference_type result = 0;
2  auto it = first.base();
3
4  // Computation when bits belong to several underlying values
5  if (first.base() != last.base()) {
6      if (first.position() != 0) {
7          result = _popcnt(*first.base() >> first.position());
8          ++it;
9      }
10     for (; it != last.base(); ++it) {
11         result += _popcnt(*it);
12     }
13     if (last.position() != 0) {
14         result += _popcnt(*last.base() << (digits - last.position()));
15     }
16     // Computation when bits belong to the same underlying value
17 } else {
18     result = _popcnt(_bextr<underlying_type>(
19         *first.base(),
20         first.position(),
21         last.position() - first.position()
22     ));
23 }
24
25 // Negates when the number of zero bits is requested
26 if (!static_cast<bool>(value)) {
27     result = std::distance(first, last) - result;
28 }
29
30 // Finalization
31 return result;
32 }
```



# Design questions

## Key design questions

- What is `bit_iterator<Iterator>::difference_type`?
- What is `bit_iterator<Iterator>::iterator_category`?
- What is `bit_iterator<Iterator>::value_type`?
- What is `bit_iterator<Iterator>::reference`?
- What is `bit_iterator<Iterator>::pointer`?

# Design questions

## Key design questions

- What is `bit_iterator<Iterator>::difference_type`?
- What is `bit_iterator<Iterator>::iterator_category`?
- What is `bit_iterator<Iterator>::value_type`?
- What is `bit_iterator<Iterator>::reference`?
- What is `bit_iterator<Iterator>::pointer`?

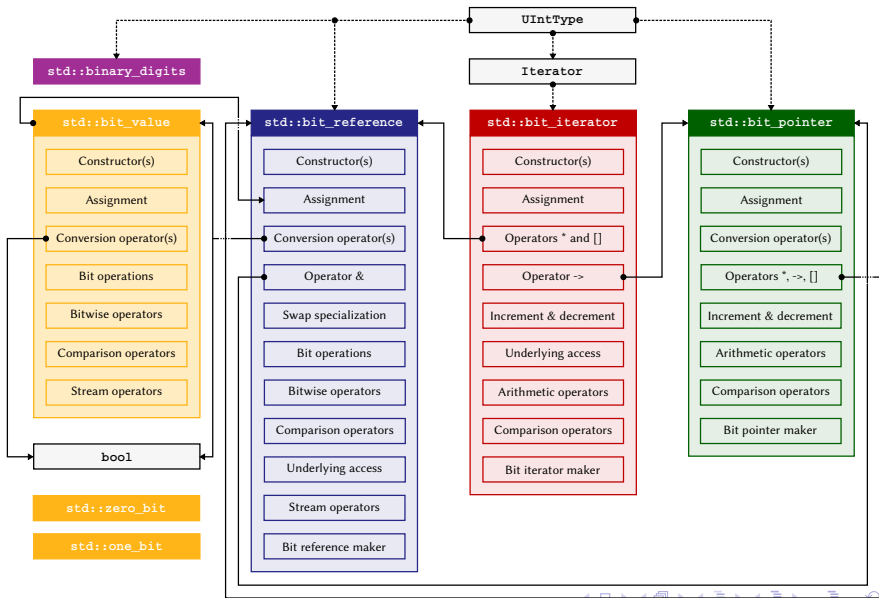
## Answer: the easy ones

- `bit_iterator<Iterator>::difference_type`  
⇒ Implementation defined, but at least as large as `std::ptrdiff_t`
- `bit_iterator<Iterator>::iterator_category`  
⇒ Same as `std::iterator_traits<Iterator>::iterator_category`

## Answer: the difficult ones

- `bit_iterator<Iterator>::value_type` ⇒ ?
- `bit_iterator<Iterator>::reference` ⇒ ?
- `bit_iterator<Iterator>::pointer` ⇒ ?

# Global architecture



## Chapter IV: On some tricky details

# What is `bit_iterator<Iterator>::value_type`?

## What is a bit?

The innocent question, that is actually very complex to answer...

## Related questions

- What functionalities a bit should provide?
- What should the arithmetic behaviour of a bit be?

## Existing problem in `std::bitset` and `std::vector<bool>`

Their boolean `value_type` has a different behaviour than their reference type leading to potential errors: as an example, the behaviour of the `operator~` is different, and `operator+=` exists for their value type, but not for their reference type.

# What is a bit?

```
1 struct field {unsigned int b : 1;};
2
3 bool b0 = false; b0 = ~b0; b0 = ~b0; // 1
4 auto x0 = std::bitset<1>{}[0]; x0 = ~x0; x0 = ~x0; // 0
5 auto f0 = field{}; f0.b = ~f0.b; f0.b = ~f0.b; // 0
6
7 bool b1 = false; b1 = ~b1; // 0
8 auto x1 = std::bitset<1>{}[0]; x1 = ~x1; // 1
9 auto f1 = field{}; f1.b = ~f1.b; // 0
10
11 bool b2 = false; b2 += 1; b2 += 1; // 1
12 auto x2 = std::bitset<1>{}[0]; x2 += 1; x2 += 1; // X
13 auto f2 = field{}; f2.b += 1; f2.b += 1; // 0
14
15 bool b3 = false; b2 = b3 + 1; b3 = b3 + 1; // 1
16 auto x3 = std::bitset<1>{}[0]; x3 = x3 + 1; x3 = x3 + 1; // 1
17 auto f3 = field{}; f3.b = f3.b + 1; f3.b = f3.b + 1; // 0
18
19 bool b4 = false; b4 += 3; // 1
20 auto x4 = std::bitset<1>{}[0]; x4 += 3; // X
21 auto f4 = field{}; f4.b += 3; // 1
```

# What is a bit?

## According to the C standard (section 3.5)

A *bit* is a unit of data storage in the execution environment large enough to hold an object that may have one of two values.

## According to the C++ standard ([intro.memory])

A *bit* is an element of a contiguous sequence forming a byte, a *byte* being the fundamental storage unit in the C++ memory model and being at least 8-bit long.

## Tentative of mathematical definition (Wikipedia)

The word *bit* stands for *binary digit*, a digit being a numeric symbol used in combinations to represent numbers in positional numeral systems.

## What is a boolean data type? (Wikipedia)

A boolean data type is a data type, having two values (usually denoted true and false), intended to represent the truth values of logic and boolean algebra.

# One of the most important slide of this talk

## Summary

- A bit is a binary digit
- A boolean is a logical data type

## A bit is not a bool

Bits and booleans are often identified, but they are two very different concepts. Both just happen to have two values.

## On `std::vector<bool>`

If bits and booleans were the same thing, `std::vector<bool>` would not raise any design issue and standardization papers about it would not exist. But they do exist...

## On `std::array<bool, N>` and `std::bitset<N>`

If bits and booleans were the same thing, `std::array<bool, N>` and `std::bitset<N>` would be equivalent. But they are not...



# What is `bit_iterator<Iterator>::value_type`? $\Rightarrow$ a `bit_value`

## class `bit_value`

Represents an independent individual bit.

## Naming: why `bit_value` and not `bit`?

Because a `bit_value` mimics a bit, but is not actually a bit since bits cannot be stored individually in memory.

## Main functionalities

- Take an unsigned integer and a bit position for construction
- No arithmetic behavior to avoid confusion (same approach as `std::byte`)
- Bitwise operators
- Flip, set and reset members
- Explicit conversion to `bool`

```
1 bit_value bval(3U, 1); // Get the bit at position 1 of 3
2 bval.flip();          // Flips the bit
3 bval = bit_value(1U); // Same as bit_value(1U, 0)
4 bval.set();           // Sets the bit to one
5 std::cout<<bval<<'\n'; // Prints the value of the bit (1)
```

# What is `bit_iterator<Iterator>::reference`? $\Rightarrow$ a `bit_reference`

```
template <class UInt> class bit_reference<UInt>
```

Represents a bit reference to a bit of an unsigned integer. Can be implemented as a reference (or a pointer) to an unsigned integer and a position.

## Main functionalities

- Same behavior as `bit_value`
- Take a reference to an unsigned integer and a bit position for construction
- Overloaded `operator&` to return a `bit_pointer`
- `address` and `position` members to get the address of the referenced unsigned integer and the position of the bit within bit

```
1 using uint_t = unsigned int;           // Sets the type of unsigned integer
2 uint_t ui = 4;                         // Creates an unsigned integer
3 bit_reference<uint_t> bref(ui, 3);      // Creates a ref to the 3rd bit of ui
4 bref.flip();                           // Flips the 3rd bit of ui
5 std::cout<<bref<<'\n';                 // Prints the 3rd bit of ui (1)
6 std::cout<<bref.position()<<'\n';     // Prints the position of the bit (3)
7 std::cout<<*(bref.address())<<'\n';   // Prints ui (12)
```

# What is `bit_iterator<Iterator>::pointer`? $\Rightarrow$ a `bit_pointer`

```
template <class UInt> class bit_pointer<UInt>
```

Represents a bit pointer to a bit of an unsigned integer. Can be implemented as a pointer to an unsigned integer and a position.

## Main functionalities

- Complementary of `bit_reference`, mimicking a pointer to a bit
- Take a pointer to an unsigned integer and a bit position for construction
- Overloaded `operator*` to return a `bit_reference`

```
1 using uint_t = unsigned int;           // Sets the type of unsigned integer
2 uint_t ui[2] = {4, 10};                // Creates an array of unsigned integer
3 bit_pointer<uint_t> bptr0(&ui[0], 3);   // Creates a pointer to the 3rd bit of ui[0]
4 bit_pointer<uint_t> bptr1(&ui[1], 8);   // Creates a pointer to the 8th bit of ui[1]
5 bptr0->flip();                          // Flips the 3rd bit of ui[0]
6 std::cout<<*bptr0<<'\n';              // Prints the 3rd bit of ui (1)
7 std::cout<<bptr0->position()<<'\n';    // Prints the position of the bit (3)
8 std::cout<<*(bptr0->address())<<'\n';  // Prints ui (12)
9 std::cout<<bptr1 - bptr0<<std::endl;   // Prints the distance in bits (37)
```

# On `binary_digits` and `bit_iterator`

```
template <class UInt> struct binary_digits
```

Helper struct inheriting from `std::integral_constant<std::size_t, N>` and giving the number of bits of unsigned integral types. Bit values, references, pointers and iterators rely on this information. Can be specialized for user types to adapt the bit library.

```
template <class Iterator> class bit_iterator
```

Combine all the preceding and provides a generic tool to manipulate bit sequences. Can be used to design bit manipulation algorithms and bit oriented data structures such as multiprecision integers.

## Main functionalities

- Based on `bit_value`, `bit_reference` and `bit_pointer`
- Take an iterator on a sequence of unsigned integers and a position for construction
- Overloaded `operator*` to return a `bit_reference`
- `base` and `position` members to get the underlying iterator and the current bit position within the current underlying unsigned integer

## Last words on bit\_iterator

```
1 // Initialization
2 using uint_t = unsigned int;
3 using container_t = std::list<uint_t>;
4 using iterator_t = typename container_t::iterator;
5 container_t container = {0, 1, 2, 3, 4};
6
7 // From the bit number 5 of container[0], to the end of the container
8 bit_iterator<iterator_t> first(std::begin(container), 5);
9 bit_iterator<iterator_t> last(std::end(container));
10
11 // Counts the bits set to 1
12 std::cout<<count(first, last, one_bit)<<std::endl;
```

### Advantages: easy to use, generic, efficient

- Very generic: can be used to reinterpret any kind of sequence of unsigned integers as a sequence of bits
- Zero overhead: most compilers can optimize the abstraction
- Acts as a standard API between users (high level) and implementers of bit manipulation algorithms or bit oriented data structures (low level)
- Good integration with the standard library: standard algorithms can be specialized to use intrinsics on bit iterators

# Chapter V: Defeating vector bool

# Preliminary words

## Previous work

Investigations have been done in the past to iterate of bit sequences efficiently. See in particular the excellent blog post by Howard Hinnant “On vector bool” (2012).

## Bit Twiddling Hacks

The webpage “Bit Twiddling Hacks” by Sean Eron Anderson has been a great source of inspiration to implement some of the bit manipulation algorithms.


## Acknowledgments

The implementation of `reverse` was done with Maghav Kumar (mkumar10@illinois.edu) at the University of Illinois at Urbana-Champaign.

# Performances: early benchmark

## Benchmark of standard algorithms on `vector<bool>` vs their `bit_iterator` specialization (logarithmic scale) [preliminary results]

Average time for 100 benchmarks with a vector size of 100,000,000 bits (speedups are provided at the top of each column)  
i7-2630QM @ 2.00GHz, Linux 3.13.0-74-generic, g++ 5.3.0, -O3, -march-native, stdlibc++ 20151204, credit: Vincent Reverdý






# Performances: early benchmark

Benchmark of standard algorithms on `vector<bool>` vs their `bit_iterator` specialization (linear scale) [preliminary results]

Average time for 100 benchmarks with a vector size of 100,000,000 bits (speedups are provided at the top of each column)  
i7-2630QM @ 2.00GHz, Linux 3.13.0-74-generic, g++ 5.3.0, -O3, -march-native, stdlibc++ 20151204, credit: Vincent Reverdý



Implementing a new version of `std::reverse` for bit iterators (1/3)

```
1 // Reverses the order of the bits in the provided range
2 template <class BidirIt>
3 void reverse(bit_iterator<BidirIt> first, bit_iterator<BidirIt> last)
4 {
5     // Assertions
6     _assert_range_viability(first, last);
7
8     // Initialization
9     using underlying_type = typename bit_iterator<BidirIt>::underlying_type;
10    using size_type = typename bit_iterator<BidirIt>::size_type;
11    constexpr size_type digits = binary_digits<underlying_type>::value;
12    const bool is_last_null = last.position() == 0;
13    size_type diff = (digits - last.position()) * !is_last_null;
14    auto it = first.base();
15    underlying_type first_value = {};
16    underlying_type last_value = {};
17
18    // Reverse when bit iterators are aligned
19    if (first.position() == 0 && last.position() == 0) {
20        std::reverse(first.base(), last.base());
21        for (; it != last.base(); ++it) {
22            *it = _bitswap(*it);
23        }
24    }
```

Implementing a new version of `std::reverse` for bit iterators (2/3)

```
1 // Reverse when bit iterators do not belong to the same underlying value
2 } else if (first.base() != last.base()) {
3     // Save first and last element
4     first_value = *first.base();
5     last_value = *std::prev(last.base(), is_last_null);
6     // Reverse the underlying sequence
7     std::reverse(first.base(), std::next(last.base(), !is_last_null));
8     // Shift the underlying sequence to the left
9     if (first.position() < diff) {
10        it = first.base();
11        diff = diff - first.position();
12        for (; it != last.base(); ++it) {
13            *it = _shld<underlying_type>(*it, *std::next(it), diff);
14        }
15        *it <<= diff;
16        it = first.base();
17        // Shift the underlying sequence to the right
18    } else if (first.position() > diff) {
19        it = std::prev(last.base(), is_last_null);
20        diff = first.position() - diff;
21        for (; it != first.base(); --it) {
22            *it = _shrd<underlying_type>(*it, *std::prev(it), diff);
23        }
24        *it >>= diff;
25        it = first.base();
26    }
27    // Bitswap every element of the underlying sequence
28    for (; it != std::next(last.base(), !is_last_null); ++it) {
29        *it = _bitswap(*it);
30    }
```

Implementing a new version of `std::reverse` for bit iterators (3/3)

```
1 // Blend bits of the first element
2 if (first.position() != 0) {
3     *first.base() = _bitblend<underlying_type>(
4         first_value,
5         *first.base(),
6         first.position(),
7         digits - first.position()
8     );
9 }
10 // Blend bits of the last element
11 if (last.position() != 0) {
12     *last.base() = _bitblend<underlying_type>(
13         *last.base(),
14         last_value,
15         last.position(),
16         digits - last.position()
17     );
18 }
19 // Reverse when bit iterators belong to the same underlying value
20 } else {
21     *it = _bitblend<underlying_type>(
22         *it,
23         _bitswap(*it >> first.position()) >> diff,
24         first.position(),
25         last.position() - first.position()
26     );
27 }
28 }
```

# Enhanced version of `std::reverse`

## Difference between the two versions

- The new version works for all cases, not only aligned bit sequences
- The fundamental low level functions like `_bitswap` have been improved (`bit_details` file for details), combining compiler intrinsics, “Bit Twiddling Hacks” and template metaprogramming



## Results

- Speed-up of the old version: 31×
- Speed-up of the new version: 86×



# Summary

## Summary: The Bit Library

- `std::vector<bool>` is broken and `std::bitset` is very limited
- `bit_iterator` is a good way to combine ease of use, genericity and performance (orders of magnitude better than `std::vector<bool>`) for bit manipulation
- Abstracting bits is not an easy task
- The Bit Library is still work in progress: specialization of most of the standard algorithms need to be done, `bit_value` and `bit_reference` functionalities are still likely to evolve
- The library is available online for everyone to test, benchmark, share and participate

## Summary: standardization

- Proposal and wording P0237 already submitted several times
- Positive feedback from LEWG
- Some issues need to be solved particularly regarding to cv-qualifiers
- We are aiming for C++Next

# Open questions and future directions

## Open questions

- Should `bit_value` be a template class to allow `bit_value&` to be implicitly convertible to `bit_reference`?
- Template and cv-qualifiers do not combine well for proxy classes such as bit abstractions: for example what should happen for a `const bit_value<volatile T>` or a `const bit_reference<T>`?  
⇒ If anyone has a clear view on that problem, please come or contact me!

## Future directions (collaborations welcome!)

- Specialization of all relevant standard algorithms for bit iterators (for high performance and low latency computing)
- Implementation of bit ranges (Range proposal)
- Implementation of container adapters to reinterpret containers as static or dynamic bitsets
- Work on multiprecision integer arithmetic



# Final words

## Credits

The drawings are coming from the webcomic `xkcd` by Randall Munroe

## Acknowledgments to the C++ community and in particular to

- Nathan Myers, Tomasz Kaminski, Lawrence Crowl, Howard Hinnant, Jens Maurer, Tony Van Eerd, Klemens Morgenstern, Vicente Botet Escriba, Odin Holmes and the other contributors of the ISO C++ Standard - Discussion and of the ISO C++ Standard - Future Proposals groups for their initial reviews and comments.
- The C++ Standards Committee
- The organizers of CPPCON 2016

## Research acknowledgments

- Vincent Reverdý and Robert J. Brunner have been supported by the National Science Foundation Grant AST-1313415. Robert J. Brunner has been supported in part by the Center for Advanced Studies at the University of Illinois.
- Vincent Reverdý is thankful to the Laboratoire Univers et Théories (LUTH), in France, where he started his early investigations on high performance trees.

# Final words



Thank you for your attention



This research has been done in/at

The Tree Building Blocks library team



The Data Science Group (LCDM), under the supervision of Prof. Robert Brunner



Department of Astronomy



University of Illinois at Urbana-Champaign (UIUC)

## Contact and links

- GITHUB: <https://github.com/vreverdy/bit>
- Contact: [vince.rev@gmail.com](mailto:vince.rev@gmail.com)
- ISO C++ proposal: P0237
- Collaborations are very welcome! Same for comments! Same for benchmarks!