

Под капотом  
стандартной библиотеки C++

# Insights into the C++ standard library

Pavel Novikov

 @cpp\_ape

R&D Align Technology

**align**

# What's ahead

- use `std::vector` by default instead of `std::list`
- `emplace_back` and `push_back`
- prefer to use `std::make_shared`
- avoid static objects with non-trivial constructors/destructors
- small string optimization
- priority queue and heap algorithms
- sorting and selection
  - guarantee of strengthened worst case complexity  $O(n \log(n))$  for `std::sort`
  - when to use `std::sort`, `std::stable_sort`, `std::partial_sort`, `std::nth_element`
- use unordered associative containers by default
- beware of missing the variable name

```
using GadgetList = std::list<std::shared_ptr<Gadget>>;

GadgetList findAllGadgets(const Widget &w) {
    GadgetList gadgets;
    for (auto &item : w.getAllItems())
        if (isGadget(item))
            gadgets.push_back(getGadget(item));
    return gadgets;
}

void processGadgets(const GadgetList &gadgets) {
    for (auto &gadget : gadgets)
        processGadget(*gadget);
}
```

```
using GadgetList = std::list<std::shared_ptr<Gadget>>;

GadgetList findAllGadgets(const Widget &w) {
    GadgetList gadgets;
    for (auto &item : w.getAllItems())
        if (isGadget(item))
            gadgets.push_back(getGadget(item));
    return gadgets;
}

void processGadgets(const GadgetList &gadgets) {
    for (auto &gadget : gadgets)
        processGadget(*gadget);
}
```



is this reasonable use of `std::list`?

# std::list



No copy/move while adding elements. 👍

Memory allocation for *each* element. 🙄

- expensive
- may fragment memory
- in general worse element locality compared to `std::vector`

# std::vector

std::vector<T,A>::push\_back() has *amortized*  $O(1)$  complexity

```
std::vector<int> v; // capacity=0
v.push_back(1); // allocate, capacity=3
v.push_back(2);
v.push_back(3);
v.push_back(4); //alloc, move 3 elements, capacity=6
v.push_back(5);
```



# std::vector

$c$  — initial capacity

$m$  — reallocation factor

To push back  $N$  elements we need

$\log_m N/c$  reallocations.

On each reallocation we move  $c \cdot m^i$  elements ( $i$  — # of reallocation)

$$N_{moves} = \sum_{i=0}^{\log_m N/c} c \cdot m^i = \frac{mN - c}{m - 1}$$

$$\frac{N_{moves}}{N} \approx \frac{1}{m - 1}$$



# std::vector



$c$  — initial capacity

$m$  — reallocation factor

To push back  $N$  elements we need


$\log_m N/c$  reallocations.

On each reallocation we move  $c \cdot m^i$  elements ( $i$  — # of reallocation)

$$N_{moves} = \sum_{i=0}^{\log_m N/c} c \cdot m^i = \frac{mN - c}{m - 1}$$

$$\frac{N_{moves}}{N} \approx \frac{1}{m - 1}$$

# std::vector



$c$  — initial capacity

$m$  — reallocation factor

To push back  $N$  elements we need

$\log_m N/c$  reallocations.

On each reallocation we move  $c \cdot m^i$  elements ( $i$  — # of reallocation)


$$N_{moves} = \sum_{i=0}^{\log_m N/c} c \cdot m^i = \frac{mN - c}{m - 1}$$

$$\frac{N_{moves}}{N} \approx \frac{1}{m - 1}$$

# std::vector

$c$  — initial capacity

$m$  — reallocation factor



To push back  $N$  elements we need

$\log_m N/c$  reallocations.

On each reallocation we move  $c \cdot m^i$  elements ( $i$  — # of reallocation)


$$N_{moves} = \sum_{i=0}^{\log_m N/c} c \cdot m^i = \frac{mN - c}{m - 1}$$

$$\frac{N_{moves}}{N} \approx \frac{1}{m - 1}$$

# std::vector

$c$  — initial capacity

$m$  — reallocation factor



To push back  $N$  elements we need

$\log_m N/c$  reallocations.

On each reallocation we move  $c \cdot m^i$  elements ( $i$  — # of reallocation)

$$N_{moves} = \sum_{i=0}^{\log_m N/c} c \cdot m^i = \frac{mN - c}{m - 1}$$

$$\frac{N_{moves}}{N} \approx \frac{1}{m - 1}$$

# std::vector

$c$  — initial capacity

$m$  — reallocation factor


To push back  $N$  elements we need

$\log_m N/c$  reallocations.

On each reallocation we move  $c \cdot m^i$  elements ( $i$  — # of reallocation)

$$N_{moves} = \sum_{i=0}^{\log_m N/c} c \cdot m^i = \frac{mN - c}{m - 1}$$

$$\frac{N_{moves}}{N} \approx \frac{1}{m - 1}$$



# std::vector

$c$  — initial capacity

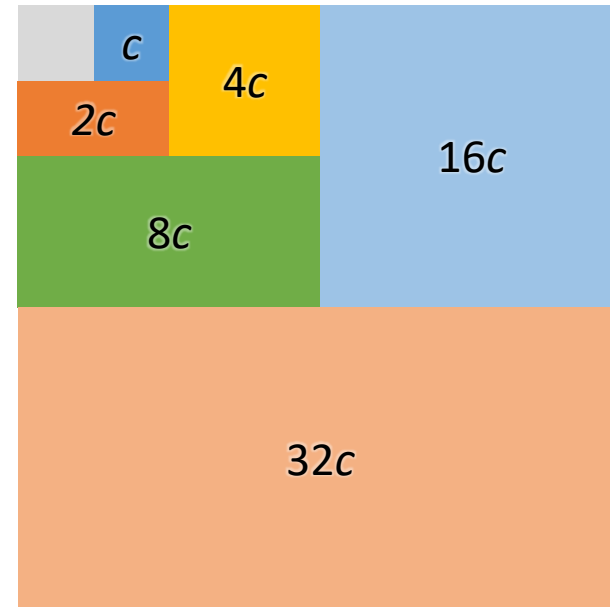
$m$  — reallocation factor

To push back  $N$  elements we need  $\log_m N/c$  reallocations.

On each reallocation we move  $c \cdot m^i$  elements ( $i$  — # of reallocation)

$$N_{moves} = \sum_{i=0}^{\log_m N/c} c \cdot m^i = \frac{mN - c}{m - 1}$$

$$\frac{N_{moves}}{N} \approx \frac{1}{m - 1}$$



# std::vector

$c$  — initial capacity


$m$  — reallocation factor

To push back  $N$  elements we need  $\log_m N/c$  reallocations.

On each reallocation we move  $c \cdot m^i$  elements ( $i$  — # of reallocation)

$$N_{moves} = \sum_{i=0}^{\log_m N/c} c \cdot m^i = \frac{mN - c}{m - 1}$$

$$\frac{N_{moves}}{N} \approx \frac{1}{m - 1}$$



# std::vector

$c$  — initial capacity


$m$  — reallocation factor

To push back  $N$  elements we need  $\log_m N/c$  reallocations.

On each reallocation we move  $c \cdot m^i$  elements ( $i$  — # of reallocation)

$$N_{moves} = \sum_{i=0}^{\log_m N/c} c \cdot m^i = \frac{mN - c}{m - 1}$$

$$\frac{N_{moves}}{N} \approx \frac{1}{m - 1}$$





`std::vector`  $N_{elements} = 64c + 1$

$c$  — initial capacity


$m$  — reallocation factor

To push back  $N$  elements we need  $\log_m N/c$  reallocations.

On each reallocation we move  $c \cdot m^i$  elements ( $i$  — # of reallocation)

$$N_{moves} = \sum_{i=0}^{\log_m N/c} c \cdot m^i = \frac{mN - c}{m - 1}$$

$$\frac{N_{moves}}{N} \approx \frac{1}{m - 1}$$



`std::vector`  $N_{elements} = 64c + 1$   
 $N_{moves} = 127c; N_{moves}/N_{elements} \approx 2$

$c$  — initial capacity

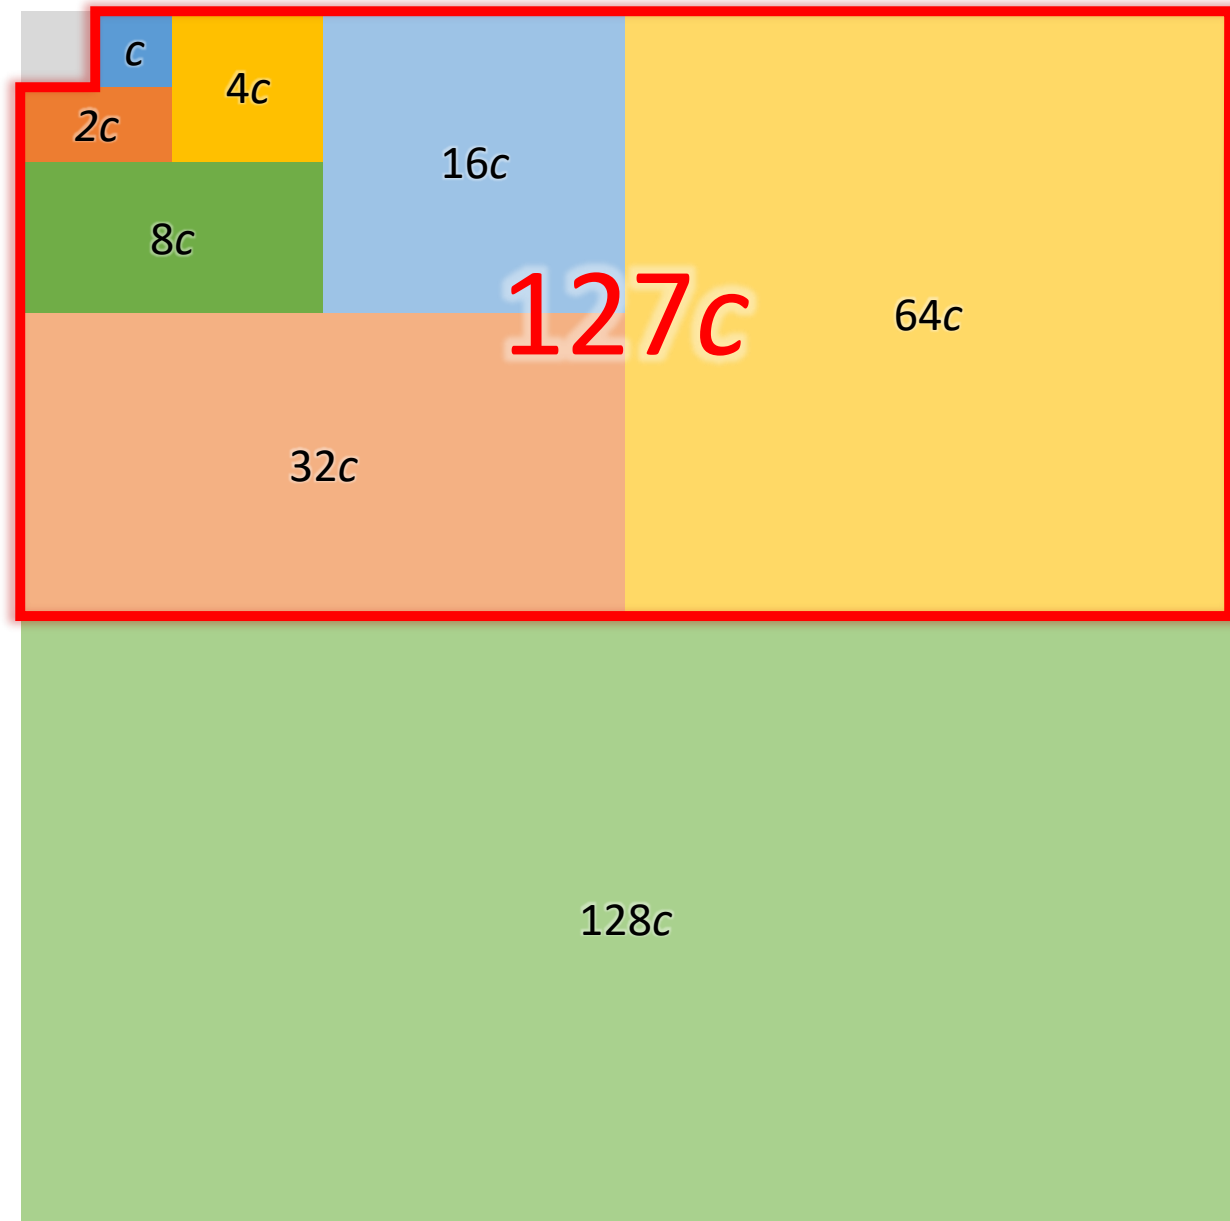
$m$  — reallocation factor

To push back  $N$  elements we need  
 $\log_m N/c$  reallocations.

On each reallocation we move  $c \cdot m^i$   
 elements ( $i$  — # of reallocation)

$$N_{moves} = \sum_{i=0}^{\log_m N/c} c \cdot m^i = \frac{mN - c}{m - 1}$$

$$\frac{N_{moves}}{N} \approx \frac{1}{m - 1}$$



# std::vector

$$N_{elements} = 64c + 1$$

$$N_{moves} = 127c; N_{moves}/N_{elements} \approx 2$$

$$N_{allocs} = 8$$

$c$  — initial capacity

$m$  — reallocation factor


To push back  $N$  elements we need

$\log_m N/c$  reallocations.

On each reallocation we move  $c \cdot m^i$  elements ( $i$  — # of reallocation)

$$N_{moves} = \sum_{i=0}^{\log_m N/c} c \cdot m^i = \frac{mN - c}{m - 1}$$

$$\frac{N_{moves}}{N} \approx \frac{m}{m - 1}$$



sum



[Browse Examples](#) [Surprise Me](#)

Assuming "sum" refers to a computation | Use as [referring to a mathematical definition](#) or a [character](#) or a [general topic](#) or a [word](#) instead

Computational Inputs:

» function to sum:

» index:

» lower limit:

» upper limit:

[Compute](#)

Sum:

$$\sum_{i=0}^{\frac{\log(N/c)}{\log(m)}} c m^i = \frac{m N - c}{m - 1}$$

[Open code](#) 

# push\_back()

```
template<typename C>
void push_back(benchmark::State &state)
{
    for (auto _ : state) {
        C container;
        for (auto counter = N; counter--;)
            container.push_back(typename C::value_type{});
        benchmark::DoNotOptimize(container);
    }
}
BENCHMARK_TEMPLATE(push_back, std::vector<T>);
BENCHMARK_TEMPLATE(push_back, std::list<T>);
```

C++

++C


`std::vector`

VS

`std::list`



# push\_back()



std::vector<int64\_t>

std::vector<BigMovable>

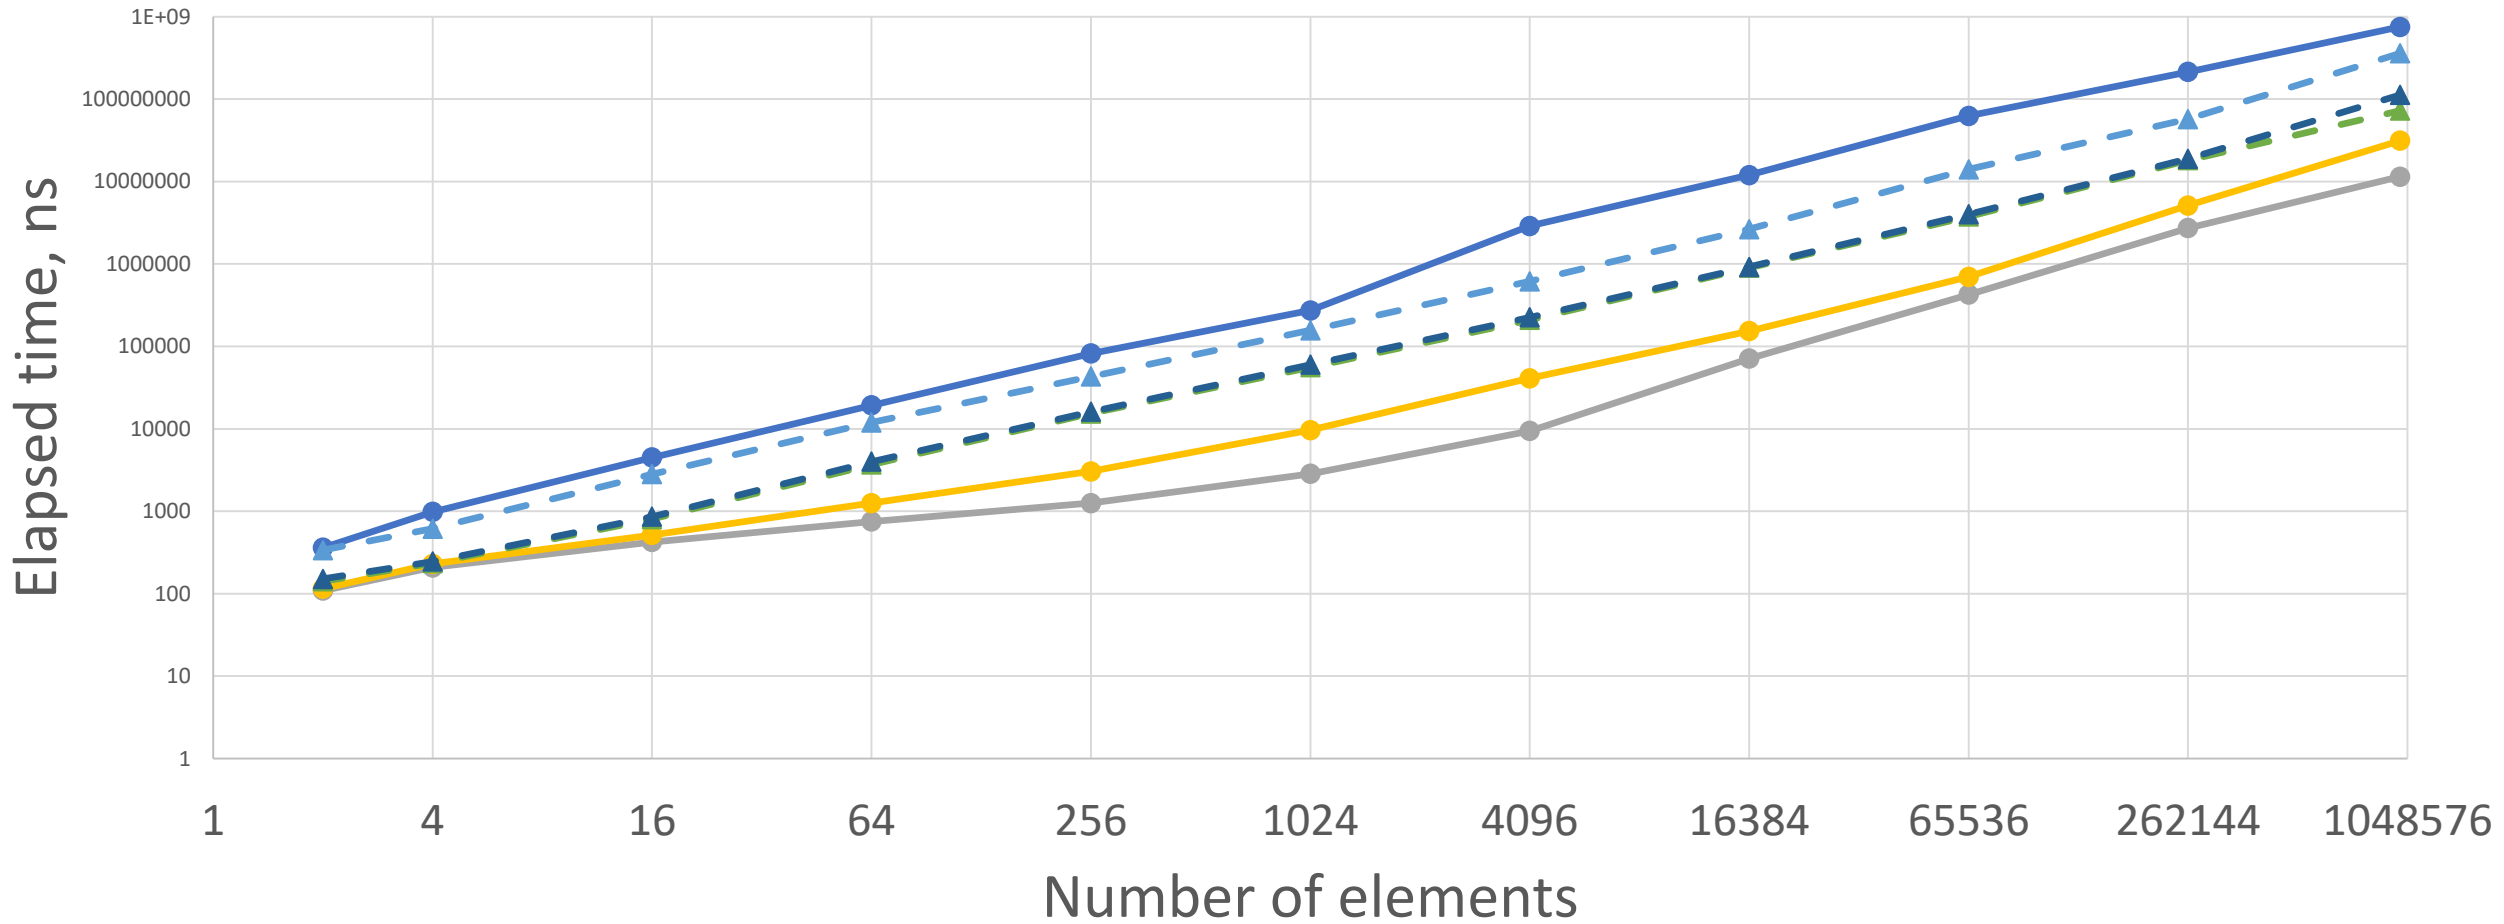
std::vector<BigCopyable>

std::list<int64\_t>

std::list<BigMovable>

std::list<BigCopyable>

# push\_back()



—●— `std::vector<int64_t>`

—●— `std::vector<BigMovable>`

—●— `std::vector<BigCopyable>`


—▲— `std::list<int64_t>`

—▲— `std::list<BigMovable>`

—▲— `std::list<BigCopyable>`




# push\_back()




# push\_back()

```
struct BigMovable {  
    BigMovable() : data{Pool::get()} {}  
    std::unique_ptr<int[], Pool::Deleter> data;  
};
```




# push\_back()



# push\_back()


```
struct BigCopyable {  
    BigCopyable() = default;  
    BigCopyable(const BigCopyable &other) :  
        data{other.data} {  
        //expensive operation  
    }  
    std::array<int, 100> data = { 1 };  
    //other data  
};
```

Elapsed time per element, ns



- std::vector<int64\_t>
- std::list<int64\_t>
- std::vector<BigMovable>
- std::list<BigMovable>
- std::vector<BigCopyable>
- std::list<BigCopyable>

# push\_back()



○ `std::vector<int64_t>`

● `std::vector<BigMovable>`

● `std::vector<BigCopyable>`

▲ `std::list<int64_t>`

▲ `std::list<BigMovable>`

▲ `std::list<BigCopyable>`

# Use `std::vector` by default

Unless

- copy and move are very expensive or unavailable;
- predominantly inserting and removing elements at random positions, splicing.

# emplace\_back and push\_back

```
std::vector<T> v;  
const auto t = T{42, true};
```

```
v.push_back(t); // copy
```

```
v.push_back(T{42, true}); // move
```

```
v.emplace_back(42, true); // construct in-place
```

```
v.emplace_back(42, true).setAwesome(true); // call method
```

```
v.emplace_back(t); // equivalent to push_back(t)
```

```
v.emplace_back(T{42, true}); // ~ push_back(T{42, true})
```

```
std::vector<T>:  
    void push_back(const T&)  
    void push_back(T&&)  
  
    template<typename... Args>  
    T &emplace_back(Args&&...)
```

# emplace\_back and push\_back

push\_back() calls only *copy* or *move* constructors.

emplace\_back() calls *any* constructor, including **explicit** constructors.

```
//was std::vector<double> values;  
std::vector<std::vector<double>> values;  
const double v = getValue();  
values.push_back(v); // does not compile  
values.emplace_back(v); // ???
```



# emplace\_back and push\_back

push\_back() calls only *copy* or *move* constructors.

emplace\_back() calls *any* constructor, including **explicit** constructors.

```
std::vector<std::unique_ptr<double>> pointers;  
pointers.push_back(&v); // does not compile  
pointers.emplace_back(&v); // ???
```

# emplace\_back and push\_back

Use emplace methods with multiple arguments.

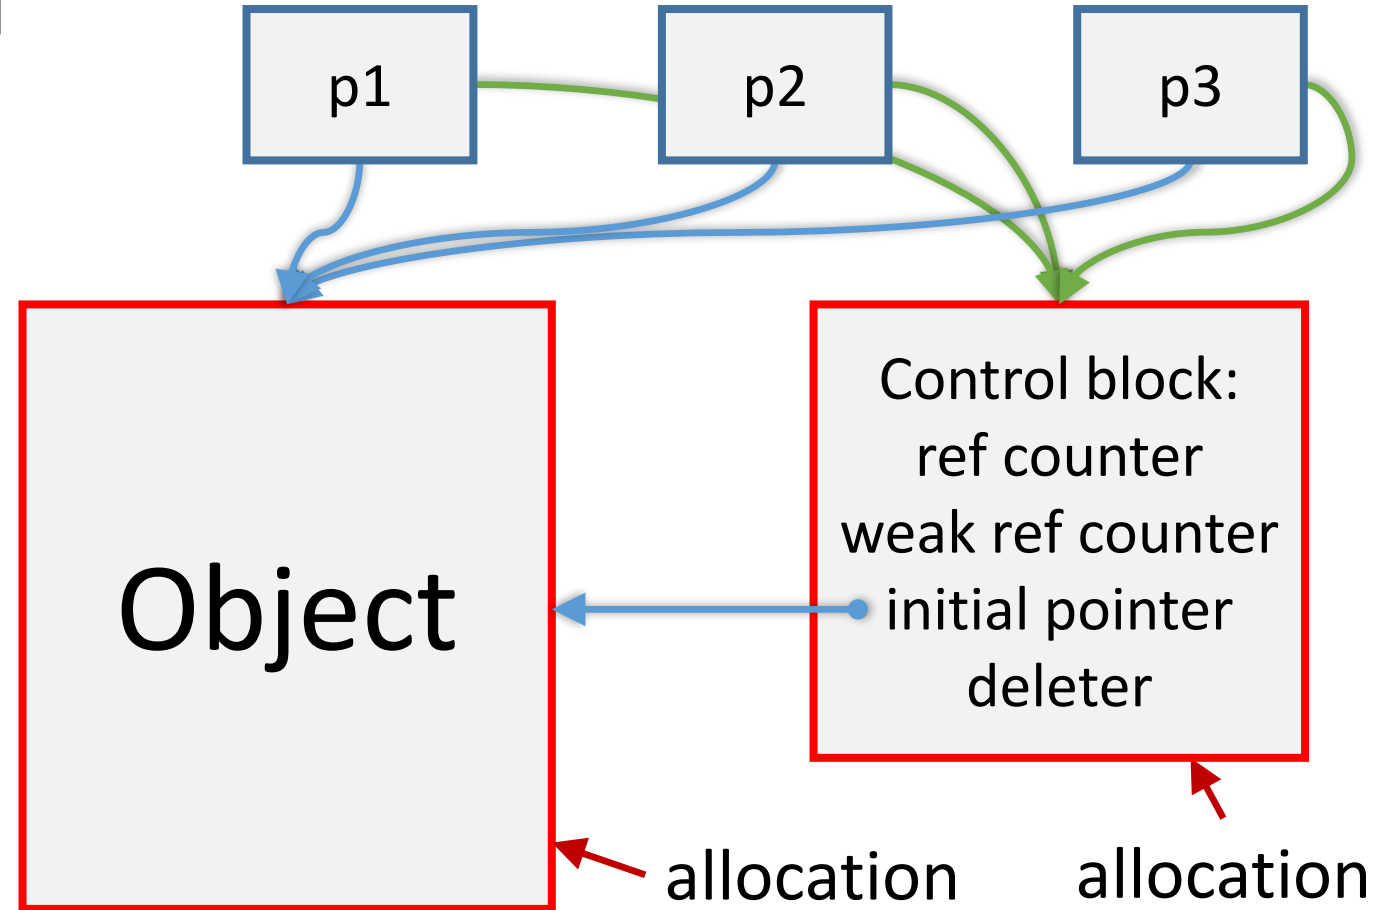
```
v.emplace_back(42, true); // construct in-place
```

Prefer push\_back() and insert() to ensure correctness when both that and emplace methods will do.

Be cautious that **emplace methods call explicit constructors.**

# std::shared\_ptr

- controls lifetime
- shares ownership

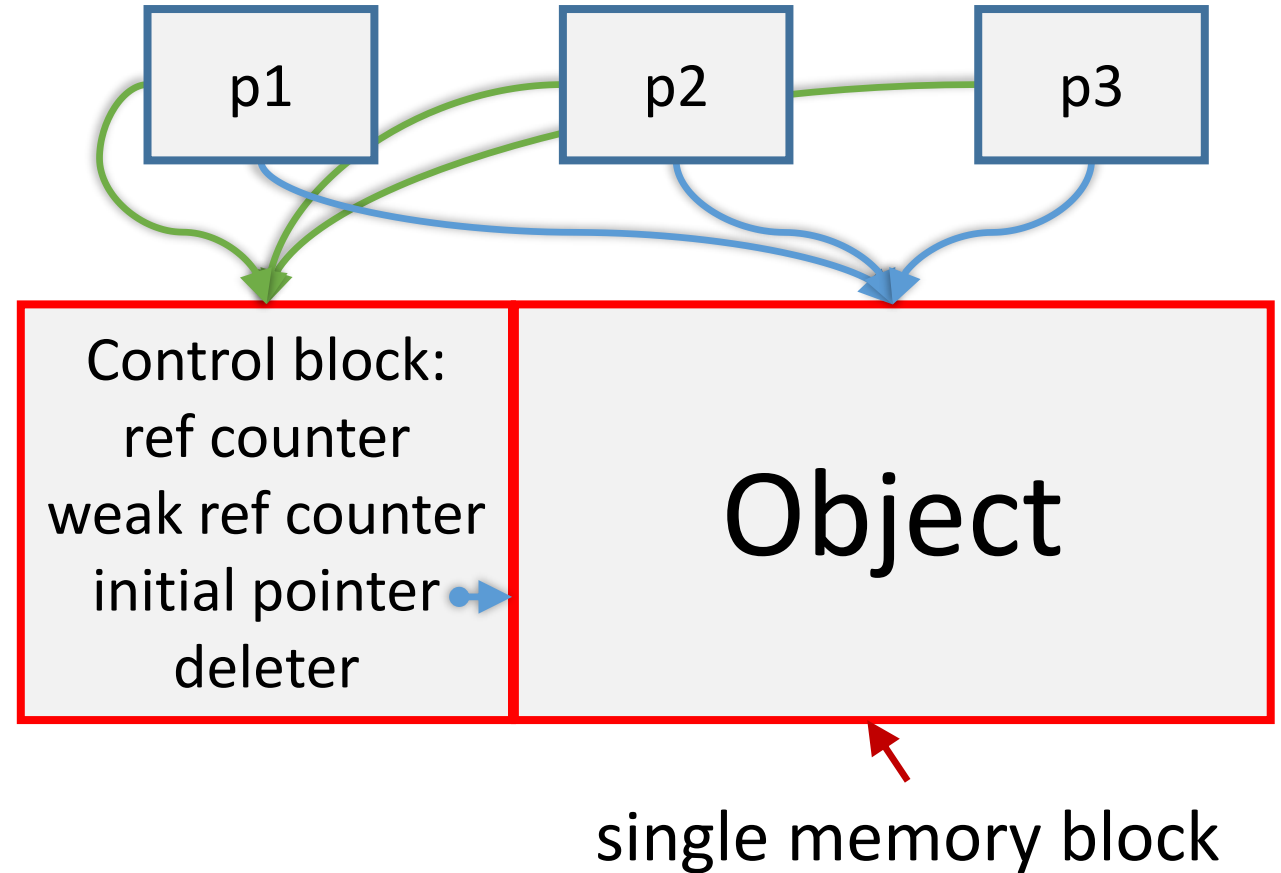


```
auto p = std::shared_ptr<Widget>(new Widget{ par });
```

# std::make\_shared

Only one memory allocation. 👍

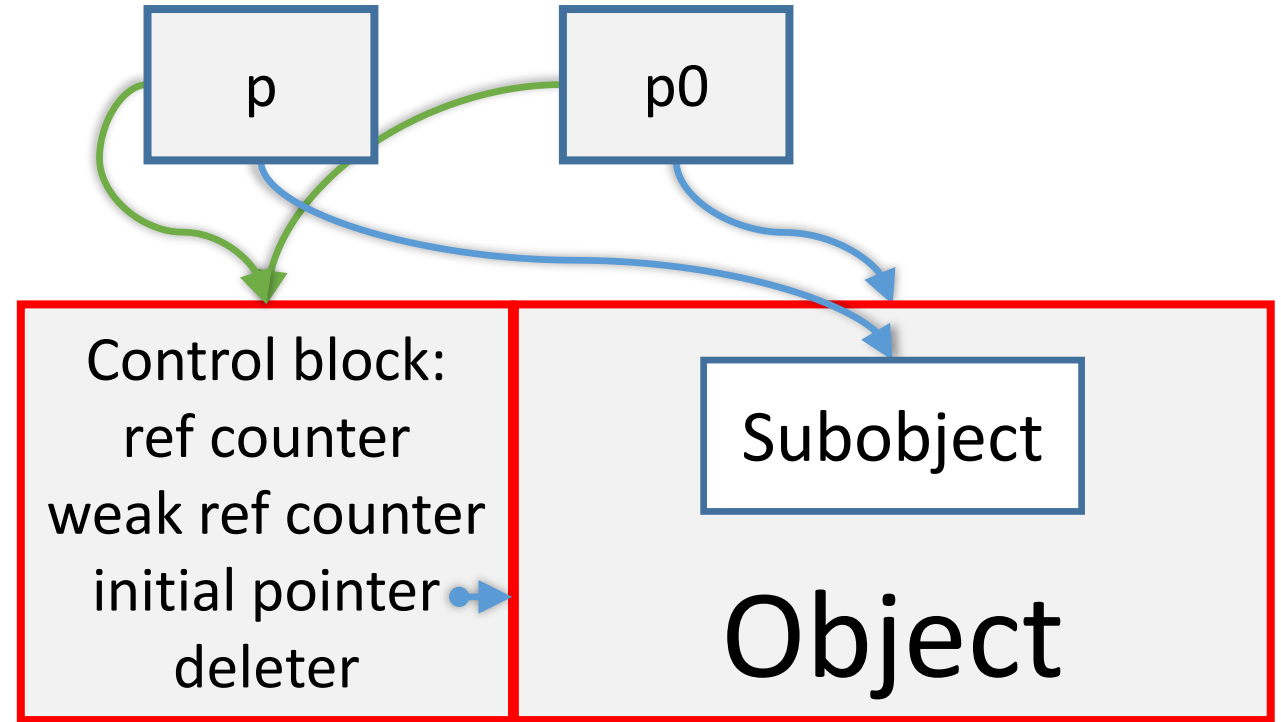
Not possible to use with a custom deleter. 🤔



```
auto p = std::make_shared<Widget>(par);
```

# Aliasing constructor

- Allows to share a part of an object  
(or anything, really)



```
auto p0 = std::make_shared<Widget>(par);  
auto p = std::shared_ptr<Gadget>(p0, &p0->gadget);
```

aliasing constructor

# Aliasing constructor

```
struct Wrapper {  
    Widget widget;  
  
    template<typename... Args>  
    Wrapper(Args&&... args) : widget{ std::forward<Args>(args)... }  
    {}  
    ~Wrapper() {  
        widget.finalize();  
    }  
};
```

```
auto p0 = std::make_shared<Wrapper>();  
auto p = std::shared_ptr<Widget>(p0, &p0->widget);
```

# Avoid static objects with non-trivial ctors/dtors

```
const std::string MsgE2BIG = "Argument list too long";
const std::string MsgEACCES = "Permission denied";
const std::string MsgEADDRINUSE = "Address in use";
...
const std::string &getMsgString(int err)
{
    switch (err)
    {
        case E2BIG: return MsgE2BIG;
        case EACCES: return MsgEACCES;
        case EADDRINUSE: return MsgEADDRINUSE;
        ...
    }
}
```

It's

 **COMPILER**  
**EXPLORER**





# Static init in MSVC

```
const std::string MsgE2BIG = "Argument list too long";
```

std::string construction

```
; void `dynamic initializer for 'MsgE2BIG'(void) PROC
```

```
sub rsp, 40 ; 00000028H
```

```
mov r8d, 22 ← strlen() elision
```

```
lea rdx, OFFSET FLAT:`string'
```

```
lea rcx, OFFSET FLAT:MsgE2BIG
```

```
call std::string::assign(char const * const,size_t)
```

```
lea rcx, OFFSET FLAT:void `dynamic atexit destructor for 'MsgE2BIG'(void)
```

```
add rsp, 40 ; 00000028H
```

```
jmp atexit ← atexit handler registration
```

# Static init in MSVC: atexit handler

```
; void `dynamic atexit destructor for 'MsgE2BIG'(void) PROC
sub rsp, 40 ; 00000028H
mov rdx, QWORD PTR MsgE2BIG+24
cmp rdx, 16
jb SHORT $LN27@dynamic
mov rcx, QWORD PTR MsgE2BIG
inc rdx
mov rax, rcx
cmp rdx, 4096 ; 00001000H
jb SHORT $LN37@dynamic
mov rcx, QWORD PTR [rcx-8]
add rdx, 39 ; 00000027H
sub rax, rcx
add rax, -8
cmp rax, 31
ja SHORT $LN55@dynamic
$LN37@dynamic:
call void operator delete(void *, size_t)
$LN27@dynamic:
movdqa xmm0, XMMWORD PTR __xmm@0000000000000000f000000000000000
movdqu XMMWORD PTR MsgE2BIG+16, xmm0
mov BYTE PTR MsgE2BIG, 0
add rsp, 40 ; 00000028H
ret 0
$LN55@dynamic:
call _invalid_parameter_noinfo_noreturn
int 3
$LN53@dynamic:
```

← inlined std::string  
destructor

# Static init in Clang

```
const std::string MsgE2BIG = "Argument list too long";
```

inlined std::string constructor

```
push rax
mov qword ptr [rip + MsgE2BIG[abi:cxx11]], offset MsgE2BIG[abi:cxx11]+16
mov qword ptr [rsp], 22
mov rsi, rsp
mov edi, offset MsgE2BIG[abi:cxx11]
xor edx, edx
call std::string::_M_create(unsigned long&, unsigned long)
mov qword ptr [rip + MsgE2BIG[abi:cxx11]], rax
mov rcx, qword ptr [rsp]
mov qword ptr [rip + MsgE2BIG[abi:cxx11]+16], rcx
movups xmm0, xmmword ptr [rip + .L.str]
movups xmmword ptr [rax], xmm0
movabs rdx, 7453016943536074612
mov qword ptr [rax + 14], rdx
mov qword ptr [rip + MsgE2BIG[abi:cxx11]+8], rcx
mov rax, qword ptr [rip + MsgE2BIG[abi:cxx11]]
mov byte ptr [rax + rcx], 0
mov edi, offset std::string::~~string() [base object destructor]
mov esi, offset MsgE2BIG[abi:cxx11]
mov edx, offset __dso_handle
call __cxa_atexit
pop rax
ret
```

strlen() elision

copying "Argument list" using SSE instructions

copying "too long" as 64 bit integer

destructor registered as atexit handler

# Avoid static objects with non-trivial ctors/dtors

```
const std::string MsgE2BIG = "Argument list too long";
const std::string MsgEACCES = "Permission denied";
const std::string MsgEADDRINUSE = "Address in use";
...
const std::string &getMsgString(int err)
{
    switch (err)
    {
        case E2BIG: return MsgE2BIG;
        case EACCES: return MsgEACCES;
        case EADDRINUSE: return MsgEADDRINUSE;
        ...
    }
}
```

# Avoid static objects with non-trivial ctors/dtors

```
const std::string_view MsgE2BIG = "Argument list too long";
const std::string_view MsgEACCES = "Permission denied";
const std::string_view MsgEADDRINUSE = "Address in use";
...
std::string_view getMsgString(int err)
{
    switch (err)
    {
        case E2BIG: return MsgE2BIG;
        case EACCES: return MsgEACCES;
        case EADDRINUSE: return MsgEADDRINUSE;
        ...
    }
}
```

# Avoid static objects with non-trivial ctors/dtors

```
const std::string_view MsgE2BIG = "Argument list too long";
const std::string_view MsgEACCES = "Permission denied";
const std::string_view MsgEADDRINUSE = "Address in use";
...
std::string_view getMsgString(int err)
{
    switch (err)
    {
        case E2BIG: return MsgE2BIG;
        case EACCES: return MsgEACCES;
        case EADDRINUSE: return MsgEADDRINUSE;
        ...
    }
}
```

# Avoid static objects with non-trivial ctors/dtors

For every static `std::string` object you waste

MSVC:  $\approx 150$  bytes of x64 instructions ( $\approx 210$  bytes in the binary)

Clang:  $\approx 100$  bytes of x64 instructions ( $\approx 160$  bytes in the binary)

To put in perspective, for  $\approx 80$  standard `errno` values  
some 12..17 KB are wasted in the binary executable

plus 8..12 KB of code in total run on *every* startup and shutdown  
of the process.

Homework: replace all static `std::strings` with  
`std::string_views` or plain arrays.  
(`strlen()` elision to the rescue!)

# Small string optimization

Naïve implementation:

```
struct NaiveString {  
    char *data = nullString;  
    size_t size = 0;  
    size_t capacity = 0;  
    static char nullString[];  
};  
char NaiveString::nullString[] = "";
```



# Small string optimization

SSO implementation:

```
struct SSOString {  
    static const size_t SSO_capacity = 16;  
    union {  
        char *data;  
        char buffer[SSO_capacity];  
    };  
    size_t size;  
    size_t capacity;  
    bool is_sso_string() const {  
        return capacity <= SSO_capacity;  
    }  
};
```

# Small string optimization

Go watch a great overview of different SSO approaches!

<https://youtu.be/kPR8h4-qZdk>



A screenshot of a presentation slide from CppCon 2016. The slide has a dark blue background with a light blue text box on the left. The text box contains the title 'std::string replacement' and two bullet points. Below the text box, it says '1% performance win'. On the right side of the slide, there is a video thumbnail showing a man on stage, with his name 'NICHOLAS ORMROD' below it. Below the thumbnail, the text reads 'The strange details of std::string at Facebook'. At the bottom right, there is a logo for 'CppCon.org'. The top right corner of the slide features the 'cppcon | 2016' logo and the text 'THE C++ CONFERENCE • BELLEVUE, WASHINGTON'.




# Heap algorithms

Largest element is at the top for *max heap*.

Smallest element is at the top for *min heap*.

Sorted vector *is also a heap*.




# Heap algorithms

Can be built in  $O(N)$  time (**std::make\_heap**).


Insertion of an element (**std::push\_heap**) and removal of the top element (**std::pop\_heap**) is  $O(\log N)$ .

Sorting (**std::sort\_heap**) is  $O(N \log N)$ .



# Heap algorithms

`std::priority_queue` is a container adaptor that wraps heap algorithms in a convenient interface.



# Sorting

State of the art sort algorithms:

**Insertion sort:**  $O(N^2)$  on average, **fastest on small sets** in practice.

**Quicksort:**  $O(N \log N)$  on average,  $O(N^2)$  in the worst case, **fastest on average** in practice.

**Heapsort:**  $O(N \log N)$  **in all cases**, slower than quicksort in practice.

# Big O notation

$$f(x) = O(g(x))$$

if and only if for some positive  $m$

$$|f(x)| \leq m \cdot g(x)$$

for all  $x \geq x_0$ .

## Properties

$f(n) = 9 \log n + 5(\log n)^4 + 3n^2 + 2n^3 = O(n^3)$  — defined by the fastest growing term.

$k \cdot f(x) = O(k \cdot g(x)) = O(g(x))$ , where  $k$  is constant

# Sorting

State of the art sort algorithms:

**Insertion sort:**  $O(N^2)$  on average, **fastest on small sets** in practice.

**Quicksort:**  $O(N \log N)$  on average,  $O(N^2)$  in the worst case, **fastest on average** in practice.

**Heapsort:**  $O(N \log N)$  **in all cases**, slower than quicksort in practice.

How to guarantee  $O(N \log N)$  complexity in the worst case?

How to employ speed of quicksort avoiding the worst case?



# Sorting

Enter **introsort**.

Best of two worlds:

- speed of quicksort on average and
- $O(N \log N)$  complexity in the worst case.

Recipe:

- do at most  $m \log N$  iterations of quicksort,
- if it's not sorted yet, do heapsort on unsorted parts.

$m$  is a small constant, say 2.

# Stable sorting


- preserves order of equivalent elements,
- complexity is  $O(N \log^2 N)$  or  $O(N \log N)$  if there is enough extra space.

`stable_sort` under the hood:

**insertion sort** of small chunks and

**merge sort** in  $O(N \log N)$  and  $O(N)$  space or

**inplace merge sort** in  $O(N \log^2 N)$  time if there isn't enough space,  
all are stable.



Can we sort faster than  $O(N \log N)$ ?

# Can we sort faster than $O(N \log N)$ ?

## Radix sort!

[https://youtu.be/zqs87a\\_7zxw](https://youtu.be/zqs87a_7zxw)




C++ now 2017  
MAY 15-20  
Aspen, Colorado, USA

### Counting Sort

- Input array:  
1 5 4 0 0
- Counting array:  
2 1 0 0 1 1 0 0
- Prefix sum:  
0 2 3 3 3 4 5 5

↖  
This Tells us where to put things



**Malte Skarupke**

Sorting in less than  
 $O(n \log n)$ :  
Generalizing and  
optimizing radix sort



cppnow.org



# Sorting and selection

`partial_sort`: “approximately”  $O(N \log k)$ .  
 $k$  first elements will be sorted.

`nth_element`:  $O(N)$  “on average”.  
Independent of  $k$ .



Should we use `partial_sort` or `nth_element` + sort all the time?

What are the use cases?

# Sorting and selection

`partial_sort`:

build a heap of the first  $k$  elements in  $O(k)$   
+ insert the rest  $(N - k)$  elements into the heap in  $O(N \log k)$   
+ sort heap in  $O(k \log k)$ ,  
which results in  $O(N \log k)$  complexity overall.

`nth_element`:

introsselect in  $O(N)$  (quickselect + heapselect to avoid worst case)

`nth_element + sort`:  $O(N) + O(k \log k)$ .

For both approaches:

$O(N)$  if  $k$  is a small constant,

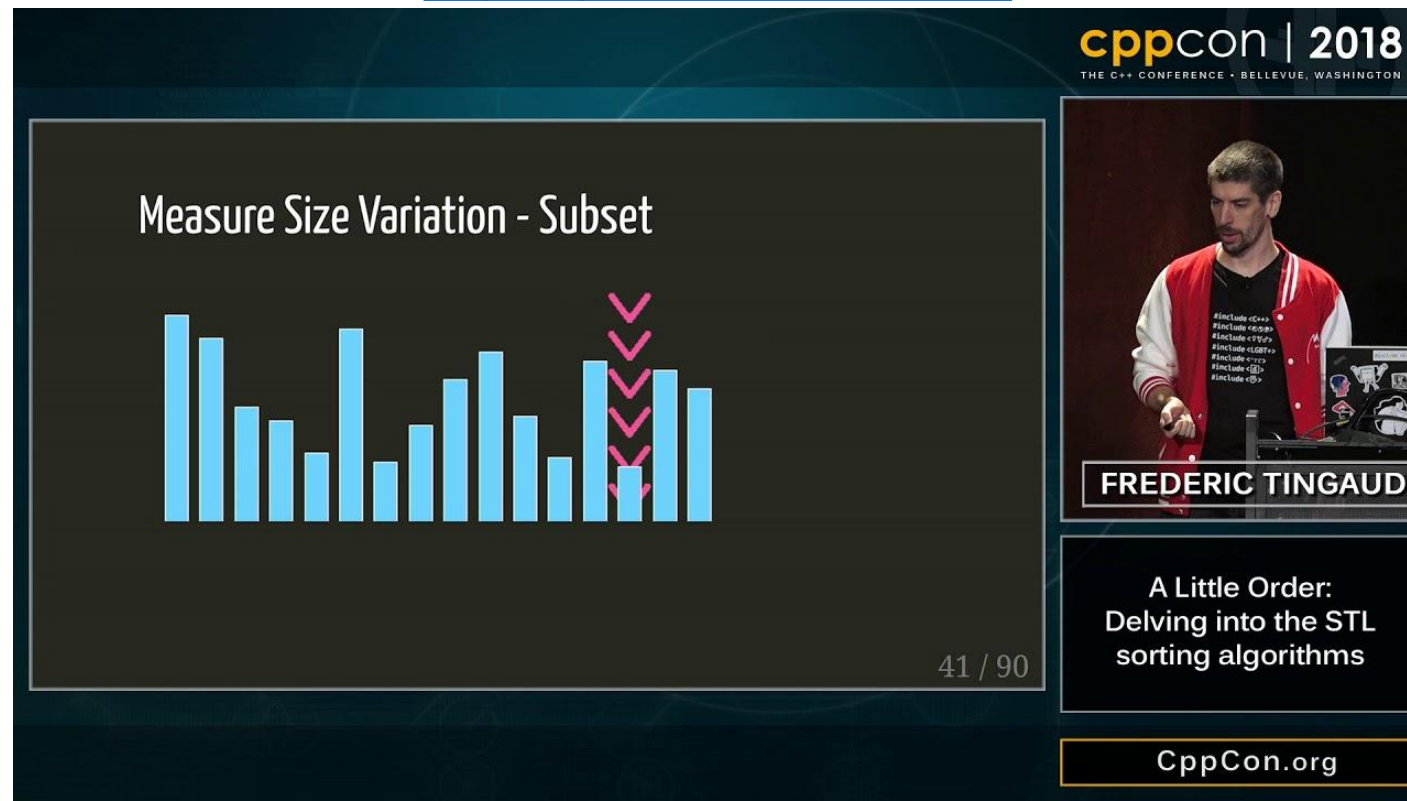
$O(N \log mN) = O(N) + O(mN \log mN) = O(N \log N)$  if  $k$  is some fraction  $mN$ .

# Sorting and selection

In practice for small constant  $k$  `partial_sort` is *faster* than `nth_element` + `sort`.

However for  $k \gg 1$  and  $N \gg 1$  `partial_sort` is significantly slower.

<https://youtu.be/-0tO3Eni2uo>



Measure Size Variation - Subset

41 / 90

cppcon | 2018  
THE C++ CONFERENCE • BELLEVUE, WASHINGTON

FREDERIC TINGAUD

A Little Order:  
Delving into the STL  
sorting algorithms

CppCon.org



# Associative containers

	insertion/deletion	lookup	
<code>unordered_map</code> <code>unordered_set</code>	$O(1)$ on average		iteration is unordered, hash function must be provided
<code>map</code> <code>set</code>	$O(\log N)$		iteration is ordered, specified by predicate
flat map/set aka sorted vector	$O(N)$	$O(\log N)$	



# flat map

```
template<typename Key, typename Value>
struct FlatMap : std::vector<std::pair<Key, Value>>
{
    Value &operator[](const Key &key);
    auto find(const Key &key) const;

private:
    struct Less;
};
```

# flat map

```
Value &operator[](const Key &key) {
    const auto i = std::lower_bound(this->begin(), this->end(),
                                    key,
                                    Less{});

    if (i == this->end()) {
        this->emplace_back(key, Value{});
        return this->back().second;
    }
    if (key < i->first)
        return this->emplace(i, key, Value{})->second;
    return i->second;
}
```

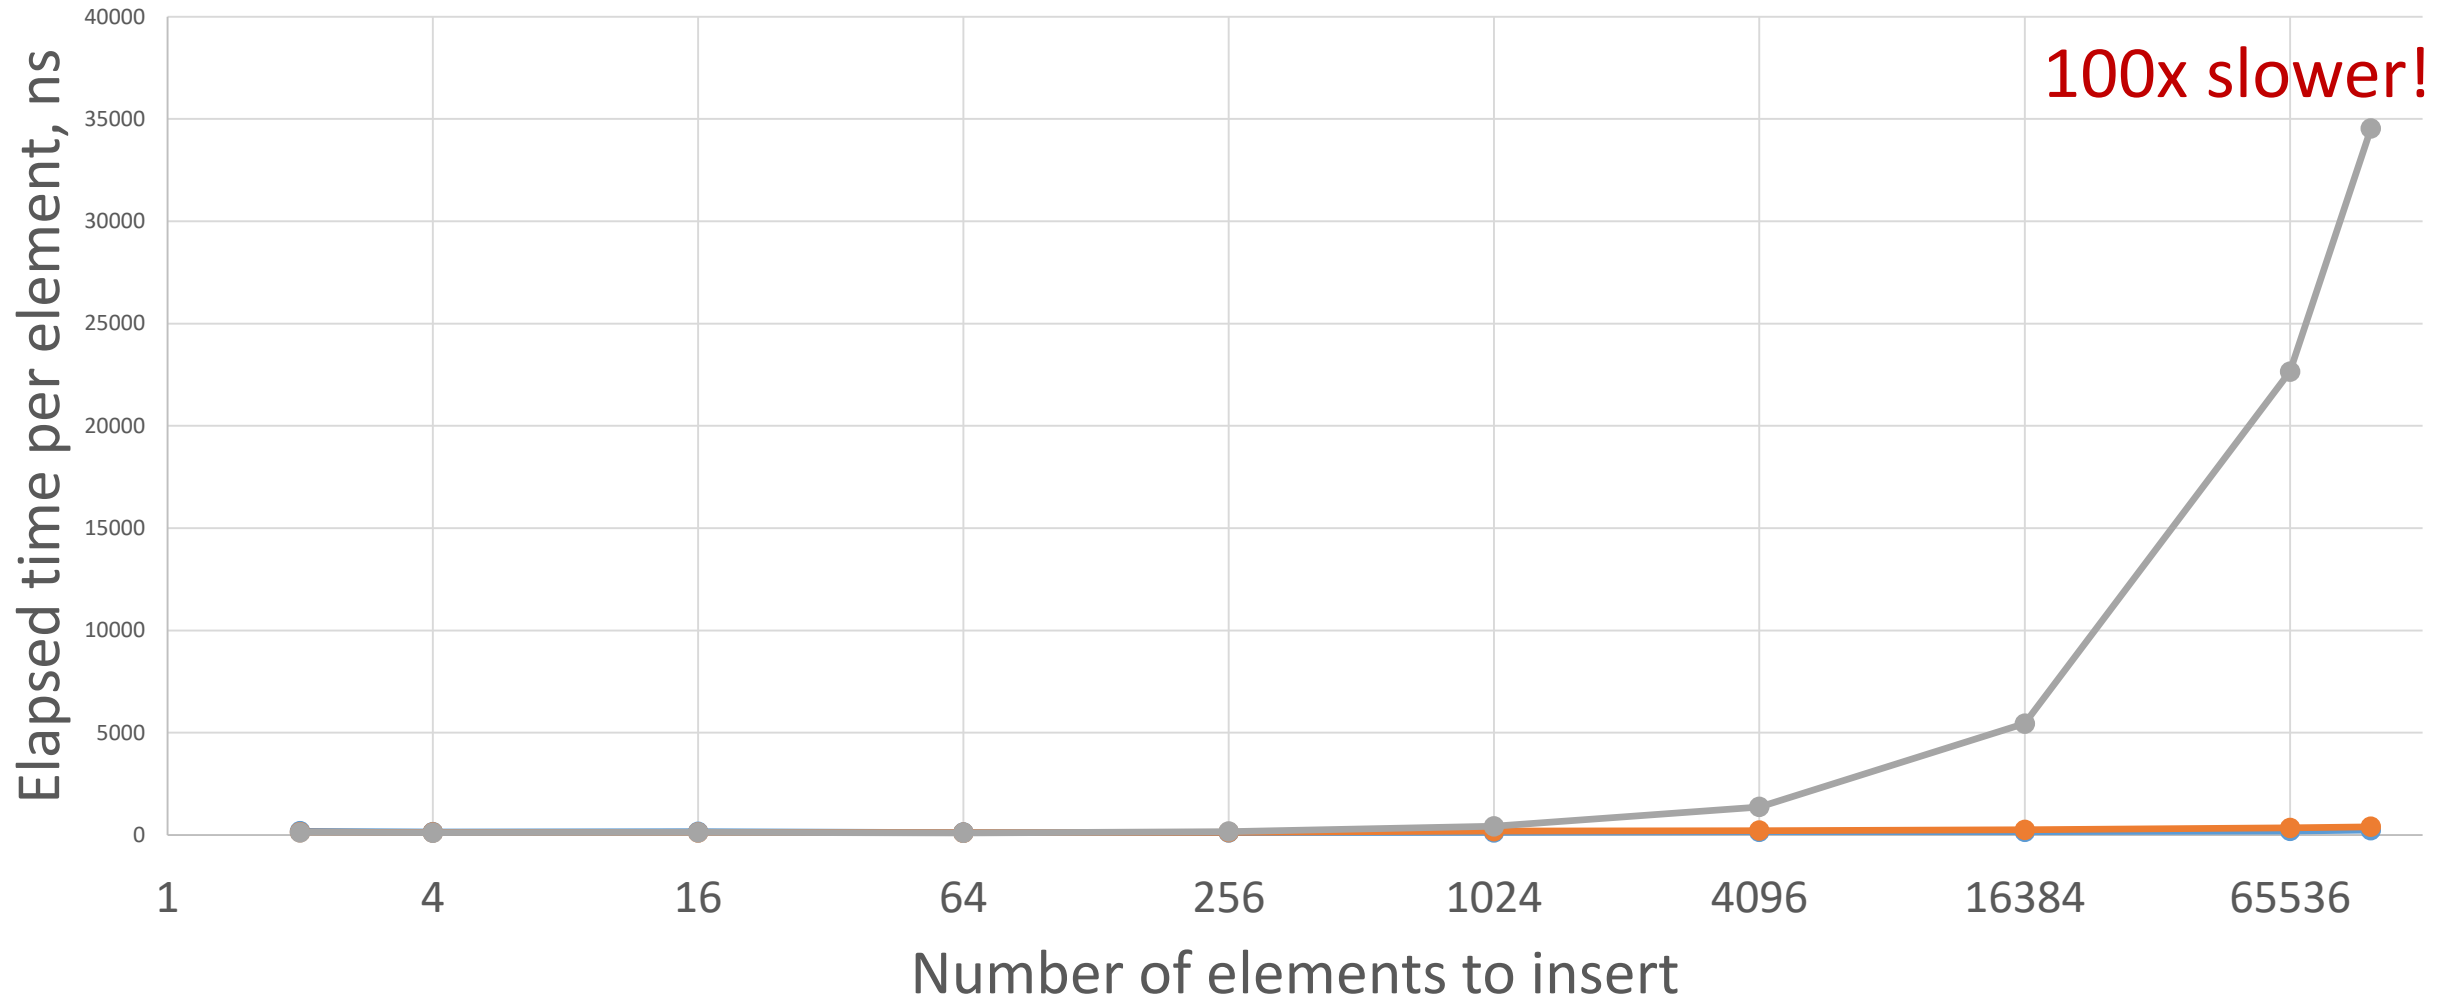
# flat map

```
auto find(const Key &key) const {  
    const auto i = std::lower_bound(this->begin(), this->end(),  
                                    key,  
                                    Less{});  
    return key < i->first ? this->end() : i;  
}
```

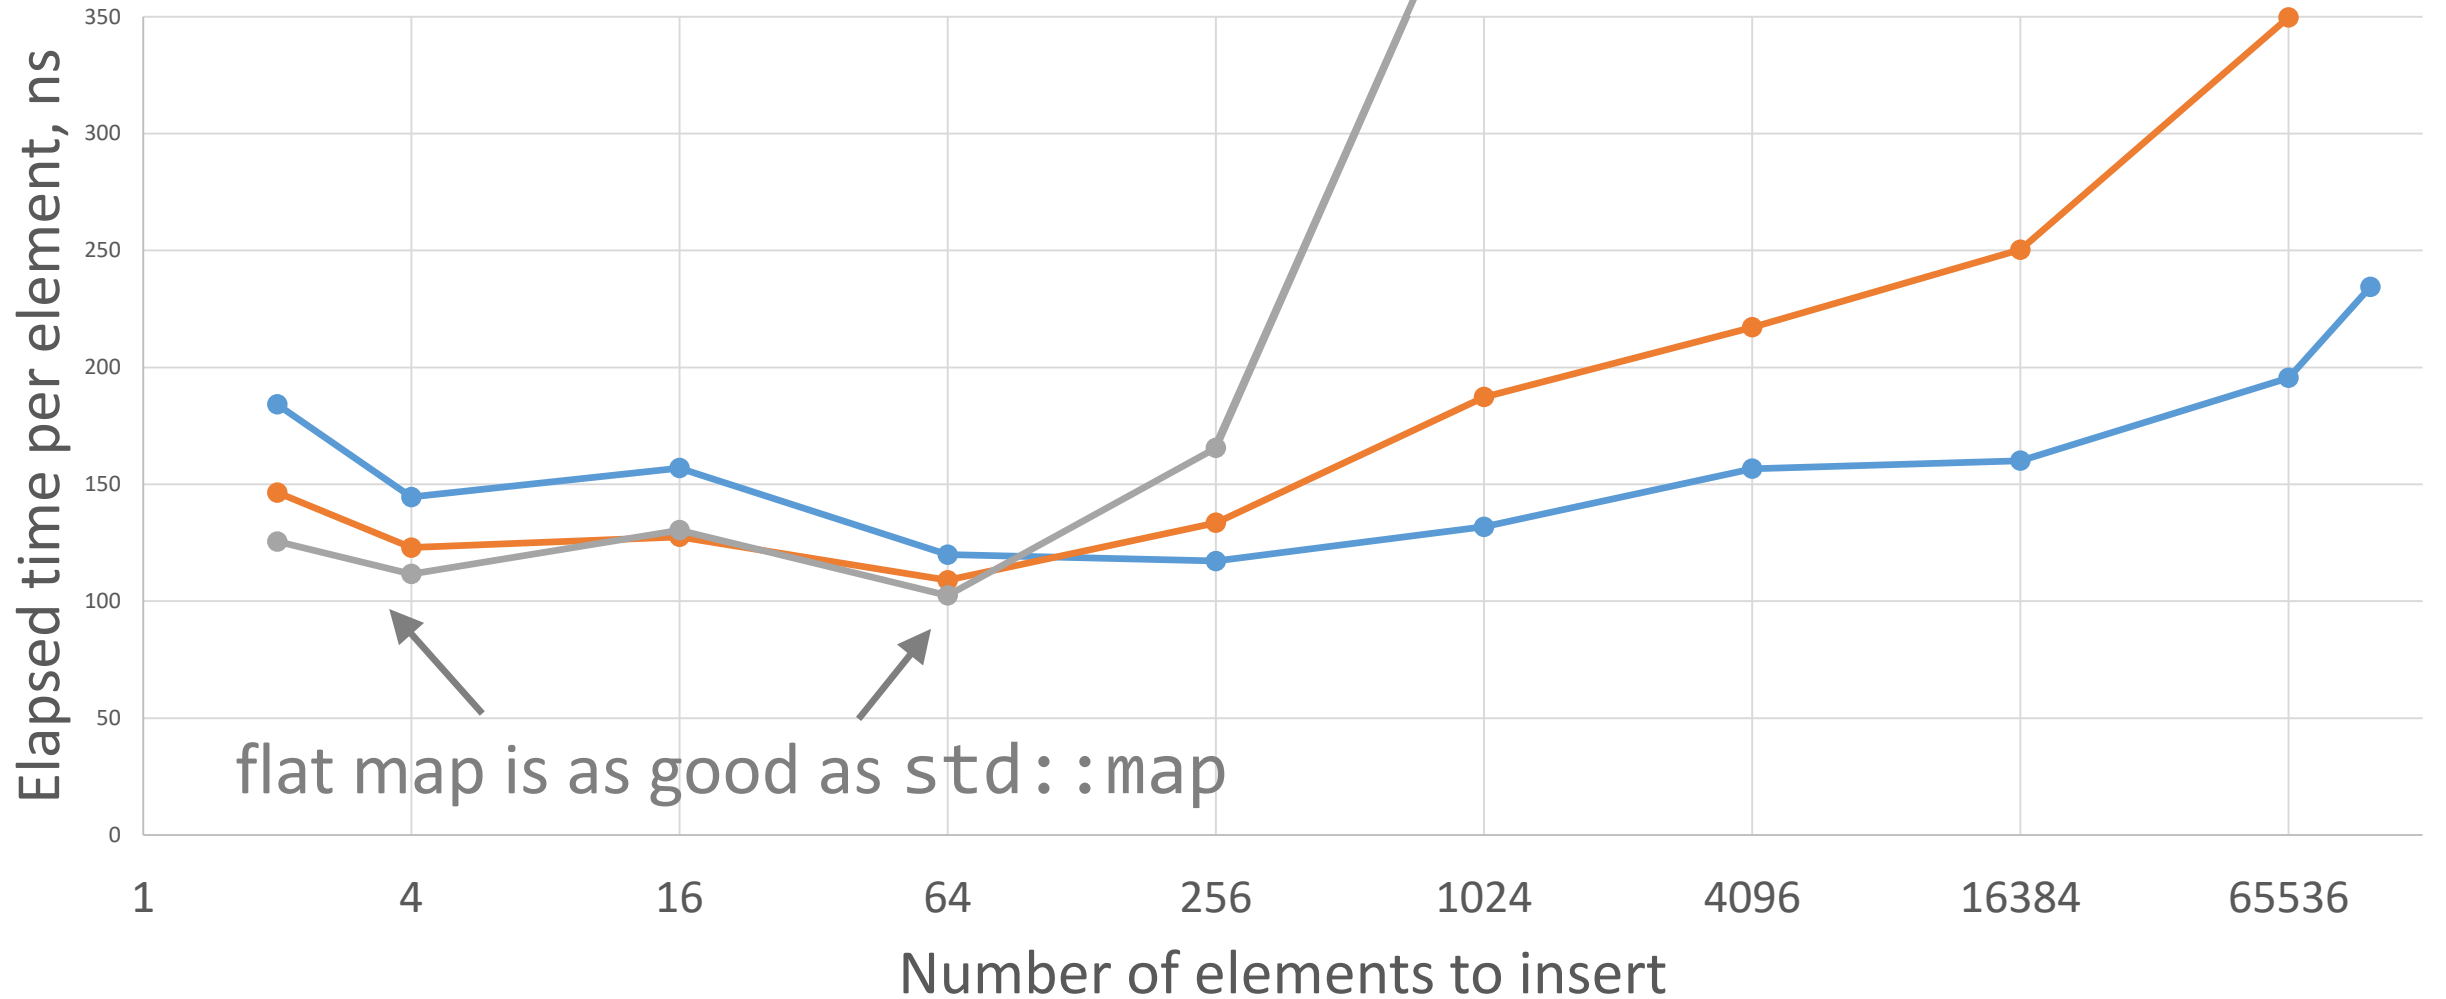
# flat map

```
struct Less
{
    template<typename A, typename B>
    bool operator()(const A &a, const B &b) {
        if constexpr (std::is_same_v<A, Key>)
            return a < b.first;
        else if constexpr (std::is_same_v<B, Key>)
            return a.first < b;
        else
            return a.first < b.first;
    }
};
```

# Insertion into map

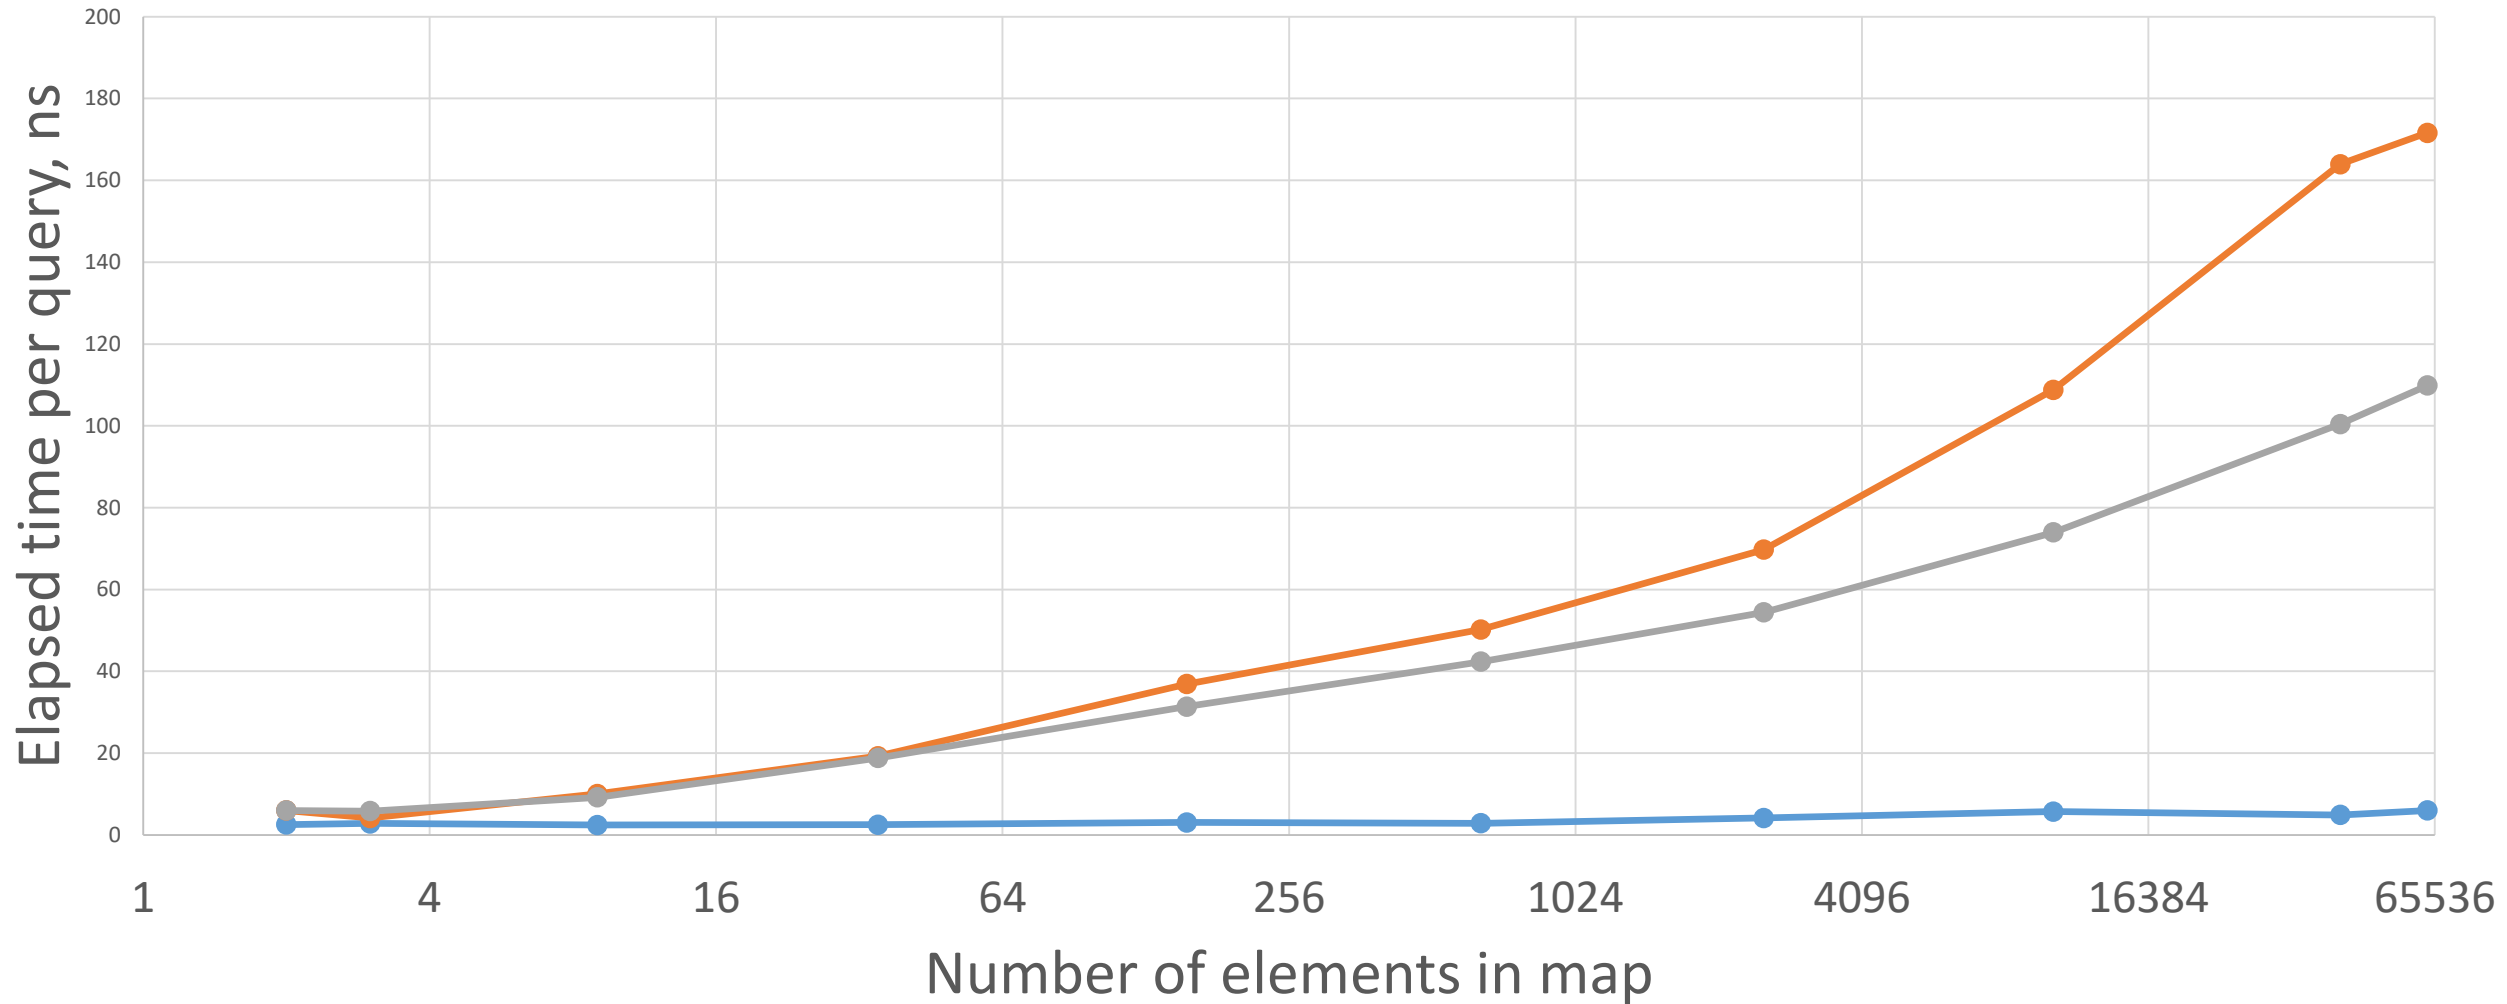


# Insertion into map



flat map is as good as `std::map`

# Querying map (Map::find)



- Use **unordered** associative containers by default
  - Use ordered associative containers if you frequently need to iterate over elements in order. Duh!
  - Use flat map only for small **constant** sizes or as constant one-time initialized map.



- Use **unordered** associative containers by default
  - Use ordered associative containers if you frequently need to iterate over elements in order. Duh!
  - Use flat map only for small **constant** sizes or as constant one-time initialized map.

What if we want to iterate over elements in order only once in a while?

# Adding elements to a map & iterating in order

Adding  $N$  elements to `std::map`:

$$O\left(\sum_{i=1}^N \log i\right) = O(N \log N - N + O(\log N)) = O(N \log N)$$

$$[u] = u + O(1)$$

Abel's summation formula

$$\begin{aligned} \sum_{i=1}^N 1 \cdot \log i &= N \log N - \int_1^N [u] \cdot (\log u)' du = N \log N - \int_1^N \frac{[u]}{u} du = \\ &= N \log N - N + 1 + O\left(\int_1^N \frac{1}{u} du\right) = N \log N - N + O(\log N) \end{aligned}$$

# Adding elements to a map & iterating in order

Adding  $N$  elements to `std::map`:  $O(N \log N)$

+ $O(N)$  iteration =  $O(N \log N)$  overall

Adding  $N$  elements to `std::unordered_map`:  $O(N)$  ( $O(1)$  per element on average)

+ $O(N)$  copying to a vector

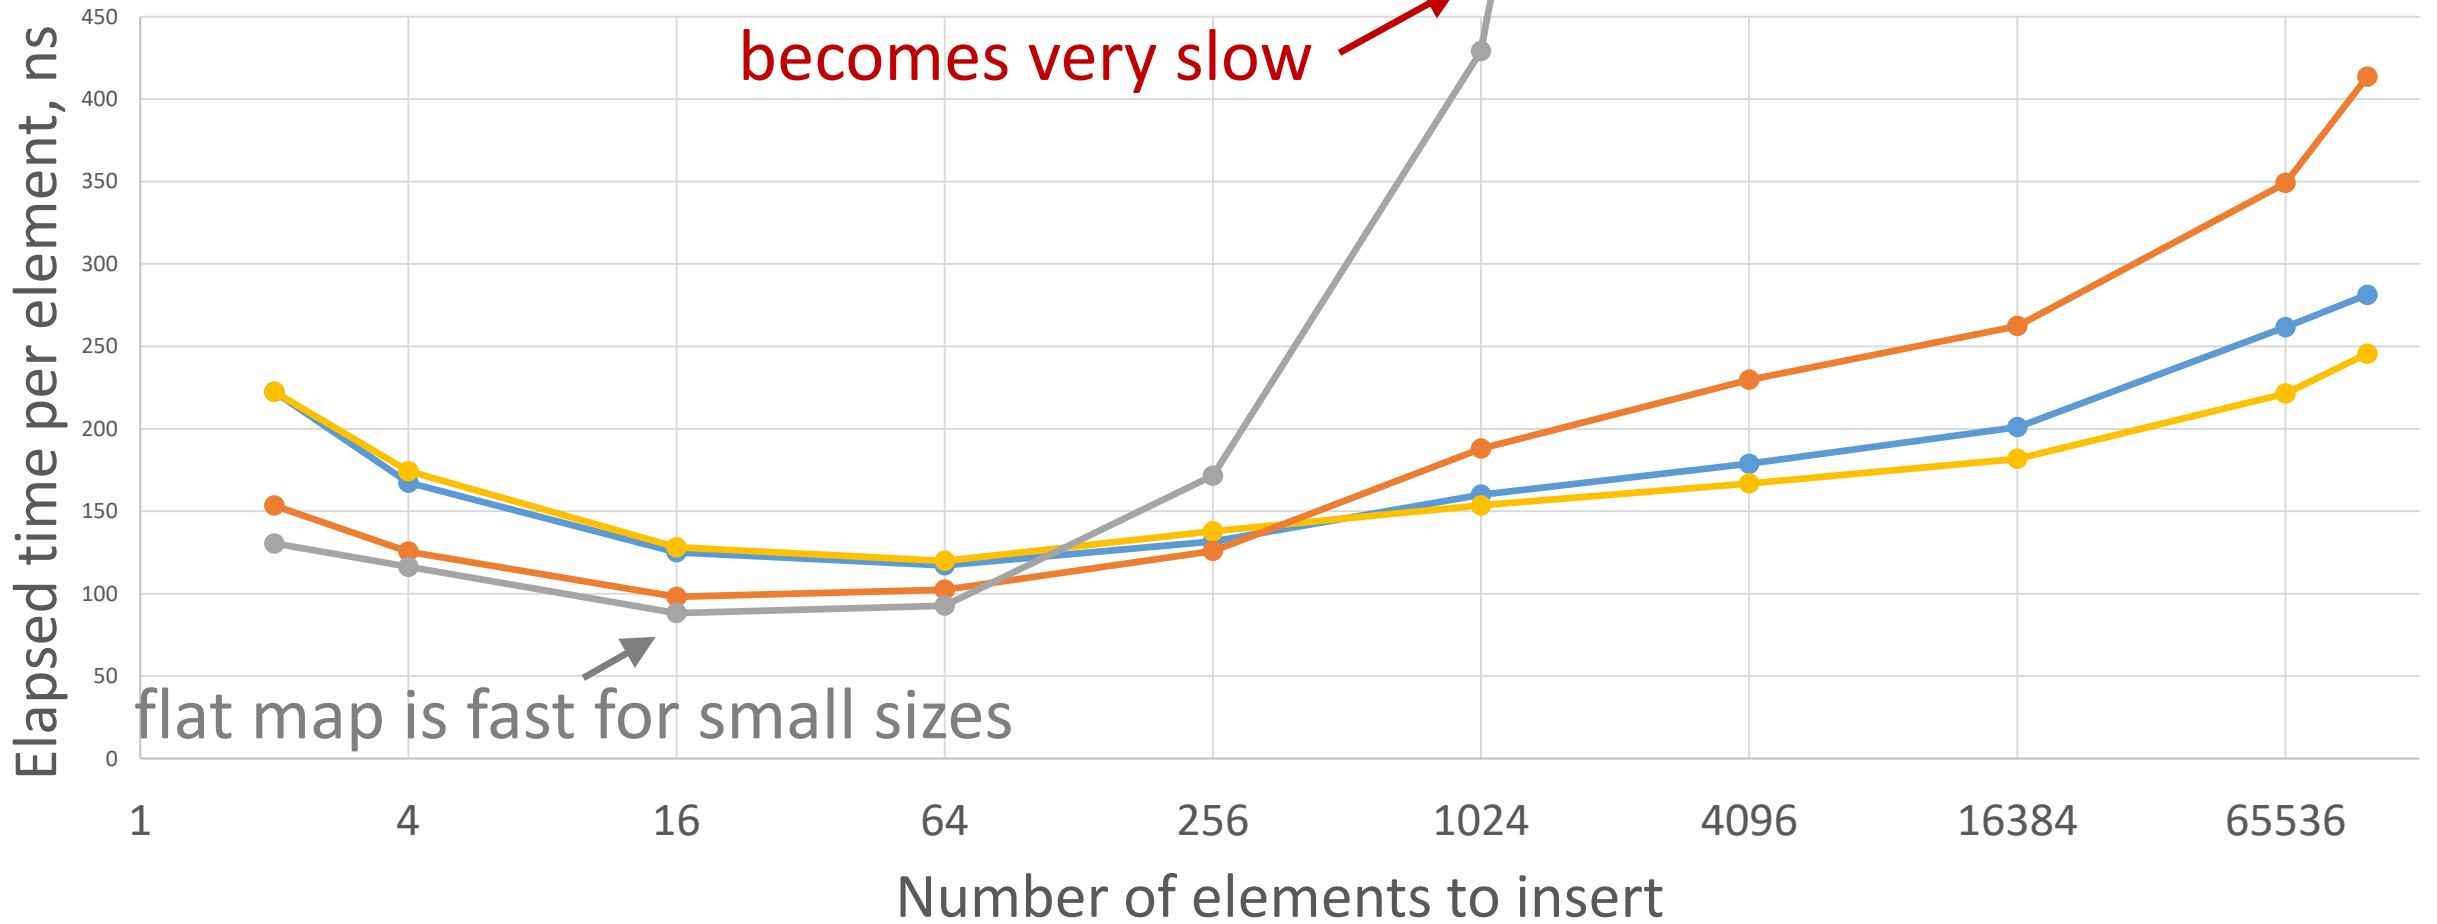
+ $O(N \log N)$  sorting

+ $O(N)$  iteration

=  $O(N \log N)$  overall

Is `std::unordered_map` approach faster than `std::map`?

# Adding elements to a map & iterating in order



— unordered\_map<int, std::string>+copy+sort

— unordered\_map<int, std::string>+copy+ska\_sort

— std::map<int, std::string>

— FlatMap<int, std::string>

# Big O notation

$$f(x) = O(g(x))$$

if and only if for some positive  $m$

$$|f(x)| \leq m \cdot g(x)$$

for all  $x \geq x_0$ .

## Properties

$f(n) = 9 \log n + 5(\log n)^4 + 3n^2 + 2n^3 = O(n^3)$  — defined by the fastest growing term.


$k \cdot f(x) = O(k \cdot g(x)) = O(g(x))$ , where  $k$  is constant

# Can hash map be faster?

Yes!

If we drop some requirements of standard library.

<https://youtu.be/M2fKMP47slQ>



C++ now 2018 MAY 7 - 11 cppnow.org

**Malte Skarupke**

You Can Do Better Than std::Unordered\_Map New and Recent Improvements to Hash Tables

Video Sponsorship Provided By:

# Pop quiz!

Does it compile?

```
std::string(someString);
```

# Pop quiz!

Does it compile?

```
std::string(someString);
```

Yes.

Does it compile?

```
std::string someString;  
std::string(someString);
```



# Pop quiz!

Does it compile?

```
std::string(someString);
```

Yes.

Does it compile?

```
std::string someString; // a variable definition  
std::string(someString); // the same
```

Redefinition of someString variable.

Pop quiz!

```
std::unique_lock<std::mutex>(mutex);
```

It compiles. See the error?

Pop quiz!

```
std::unique_lock<std::mutex>(mutex);
```

It compiles. See the error?

<https://youtu.be/lkgszkPnV8g>

A screenshot of a presentation slide from CppCon 2017. The slide has a black background with the word "mitigation" in white text at the top. Below it is a bulleted list with three items: "none", "const correctness?", and "ban it?". At the bottom of the slide is a yellow laughing emoji holding a sign that says "LOL". The slide is part of a video recording, as indicated by the "cppcon | 2017" logo and the speaker's name "LOUIS BRANDY" in the bottom right corner. The video title "Curiously Recurring C++ Bugs at Facebook" and the website "CppCon.org" are also visible.

cppcon | 2017  
THE C++ CONFERENCE • BELLEVUE, WASHINGTON

# mitigation

- none
- const correctness?
- ban it?

LOUIS BRANDY

Curiously Recurring C++ Bugs at Facebook

CppCon.org



Thanks for coming!

# Insights into the C++ standard library

Pavel Novikov

 @cpp\_ape

R&D Align Technology

# align

Slides: <https://git.io/Je44v>



# references

- `emplace_back` and `push_back`

<https://stackoverflow.com/a/36919571/11068024>: Why would I ever use `push_back` instead of `emplace_back`?

<https://abseil.io/tips/112>: Abseil Tip of the Week #112: `emplace` vs. `push_back`

- Small string optimization

<https://youtu.be/kPR8h4-qZdk>: Nicholas Ormrod “The strange details of `std::string` at Facebook”

- Sorting and selection

<https://youtu.be/-0tO3Eni2uo>: Fred Tingaud “A Little Order: Delving into the STL sorting algorithms”

[https://youtu.be/zqs87a\\_7zxw](https://youtu.be/zqs87a_7zxw): Malte Skarupke “Sorting in less than  $O(n \log n)$ : Generalizing and optimizing radix sort”

- Associative containers

<https://math.stackexchange.com/a/1560721/702319>: showing that the partial sums of  $\log(j) = n \log(n) - n + O(\log(n))$

<https://youtu.be/M2fKMP47sIQ>: Malte Skarupke “You Can Do Better than `std::unordered_map`: New Improvements to Hash Table Performance”

# references

- Beware of missing the variable name

<https://youtu.be/lkgszkPnV8g>: Louis Brandy “Curiously Recurring C++ Bugs at Facebook”