


C++ in the Audio Industry, Episode III

Align, Vectorise, Cache, Jump!



Timur Doumler

C++Now 2016, 13 May 2016

(c) 2016 by Timur Doumler

Previous Episodes

- C++ in the Audio Industry
CppCon 2015
- C++ in the Audio Industry, Episode II: Floating atomics
JUICE Summit 2015

Background



Visual Studio 2008-15
Xcode (OSX + iOS)
make (Linux)
Android Studio
gradle
ant
Live-coding

Graphics

GUI Widgets

Network

...

Audio

Build system



iOS

The screenshot displays the Traktion music production software interface. At the top, there are five tracks, each with a list of audio samples and a waveform. The tracks are labeled with sample names such as "Bell Steppa 175bpm.wav", "300b 175bpm.wav", "A-Men 150bpm.wav", "Bell 175bpm.wav", "4 Bar Roll 170bpm.wav", "Bell Steppa 175bpm.wav", "Cape 175bpm.wav", "84 Break 170bpm.wav", "Blue Shirt 175bpm.wav", "Cape 175bpm.wav", "Cape In Hats 175bpm.wav", "But Picked 150bpm.wav", "Crisis 1 170bpm.wav", "Cape In Hats 175bpm.wav", "Cape X-Tra 175bpm.wav", "Control Freak 175bpm.wav", "Crisis 2 175bpm.wav", "Cops In Hats 175bpm.wav", "Dollar Man K+S 175bpm.wav", "Catch 150bpm.wav", "E2 Break 175bpm.wav", "Cops X-Tra 175bpm.wav", "Giant Leaps 175bpm.wav", "Drop Out 175bpm.wav", "Hot Pants 150bpm.wav", "Graftid 1 175bpm.wav", "Giant Leaps 175bpm.wav", "Heart 175bpm.wav", "Head Sort 175bpm.wav", "Ringo 150bpm.wav", "Graftid 2 175bpm.wav", "Heart 175bpm.wav", "House Souther 175bpm.wav", "Scotty 150bpm.wav", "In Touch 175bpm.wav", "Ride Like 175bpm.wav", "Tomo Jamba 175bpm.wav", "What A Filler 170bpm.wav", "300b 175bpm.wav", "A-Men 150bpm.wav", "Bell 175bpm.wav", "4 Bar Roll 170bpm.wav", "Bell Steppa 175bpm.wav", "Cape 175bpm.wav", "Cape In Hats 175bpm.wav", "Cape X-Tra 175bpm.wav", "Dollar Man K+S 175bpm.wav", "Giant Leaps 175bpm.wav", "Heart 175bpm.wav", "House Souther 175bpm.wav", "Ringo 150bpm.wav", "Scotty 150bpm.wav", "In Touch 175bpm.wav", "Ride Like 175bpm.wav", "Tomo Jamba 175bpm.wav", "What A Filler 170bpm.wav".

The middle section shows three tracks: Track 1 (green) with a waveform, Track 2 (blue) with a waveform, and Track 3 (green) with a MIDI piano roll. The bottom section features a menu, a MIDI piano roll with a keyboard, and a mixer panel with various controls and a plugin list.

Traktion

NOISE



Safe & lock-free thread synchronisation.

thread reading
data from disk

audio thread



audioCallback

audioCallback

audioCallback

**REAL-TIME THREAD
DO NOT BLOCK**

audioCallback


GUI thread



messageLoop

messageLoop

thread getting
input from keyboard



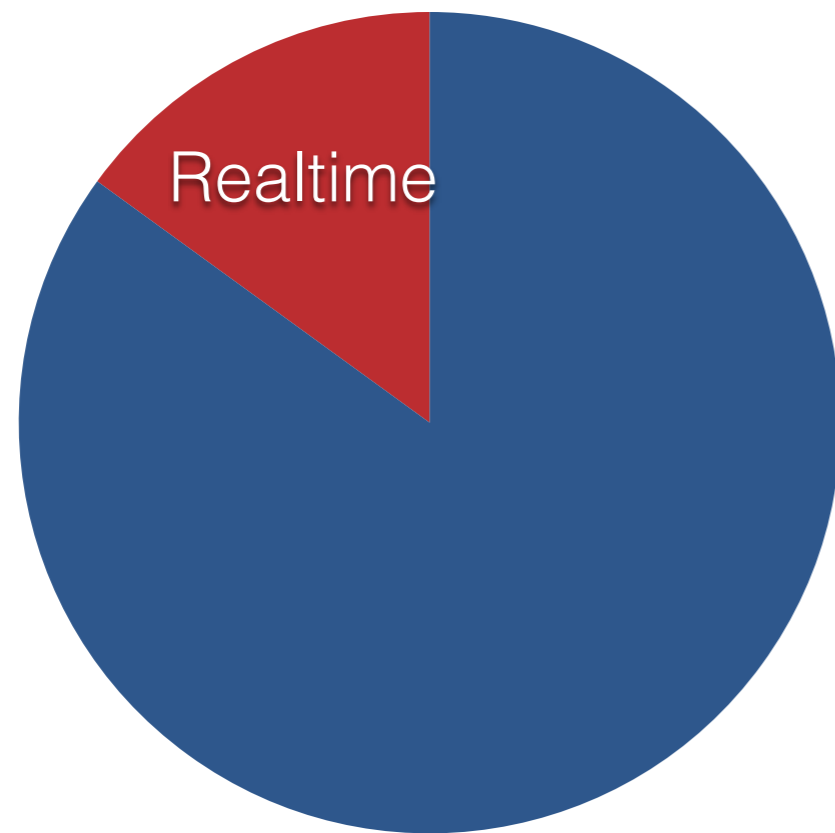
Realtime audio callback



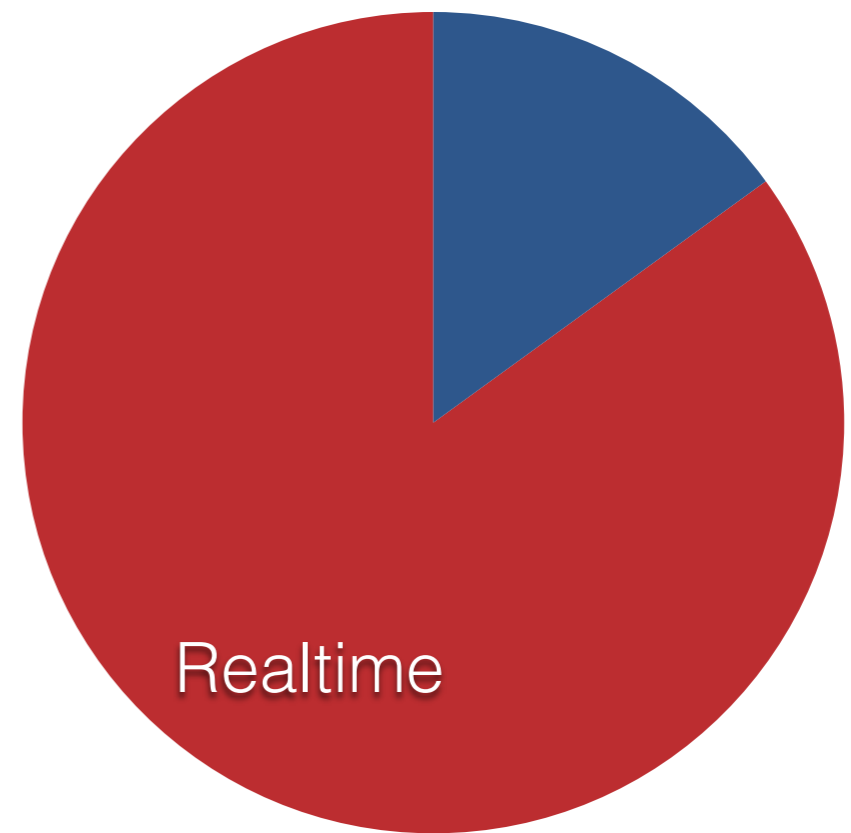
```
void audioCallback (float** channelData,  
                    int numChannels,  
                    int numSamples)  
{  
    for (int channel = 0; channel < numChannels; ++channel)  
        for (int sample = 0; sample < numSamples; ++sample)  
            channelData[channel][sample] = ...  
}
```

this should be as
fast as possible

Amount of code



Computational time





cache locality

cache
associativity

memory
alignment

(false)
sharing

vectorisation
& SIMD

branch
prediction

denormals



Devices used for benchmarks



i7



Core2Duo



iPhone6S

AppleA9



Nexus5X

Snapdragon 808

Compilers used for benchmarks

clang (Xcode 7.3 / Android Studio 2.1)

```
clang++ -std=c++11 -stdlib=libc++ -O3
```

GCC 5.3.0 (Homebrew)


```
g++-5 -std=c++11 -O3
```

Visual Studio 2015


```
cl.exe /Ox
```


2D Array traversal

```
int array[n][m];
```



List (newly allocated)



2D Array traversal


```
int array[n][n];
```

```
BENCHMARK_START(rowMajor, 100)
{
    for (int i = 0; i < n; ++i)
        for (int j = 0; j < n; ++j)
            array[i][j] += j;
}
BENCHMARK_STOP(rowMajor)
```

```
BENCHMARK_START(columnMajor, 100)
{
    for (int i = 0; i < n; ++i)
        for (int j = 0; j < n; ++j)
            array[j][i] += j;
}
BENCHMARK_STOP(columnMajor)
```

2D Array traversal benchmark

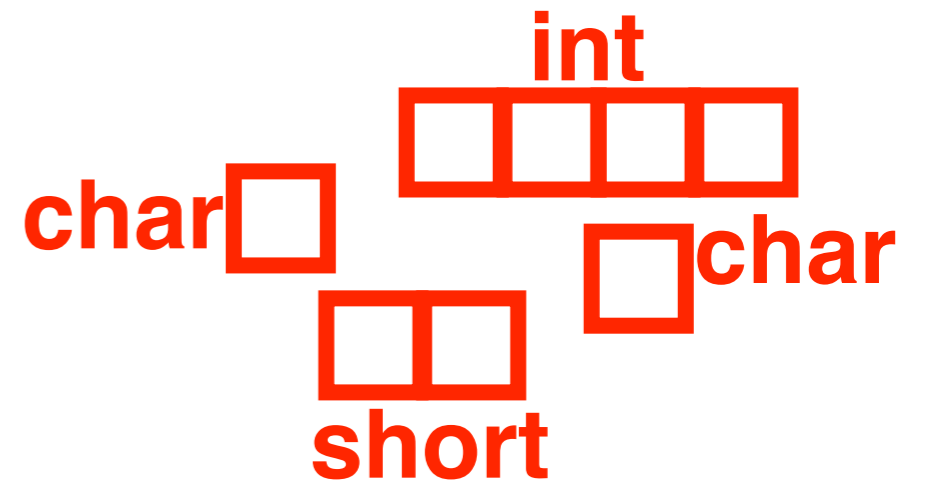
(36 MB array)



Memory



Data



Registers

CPU

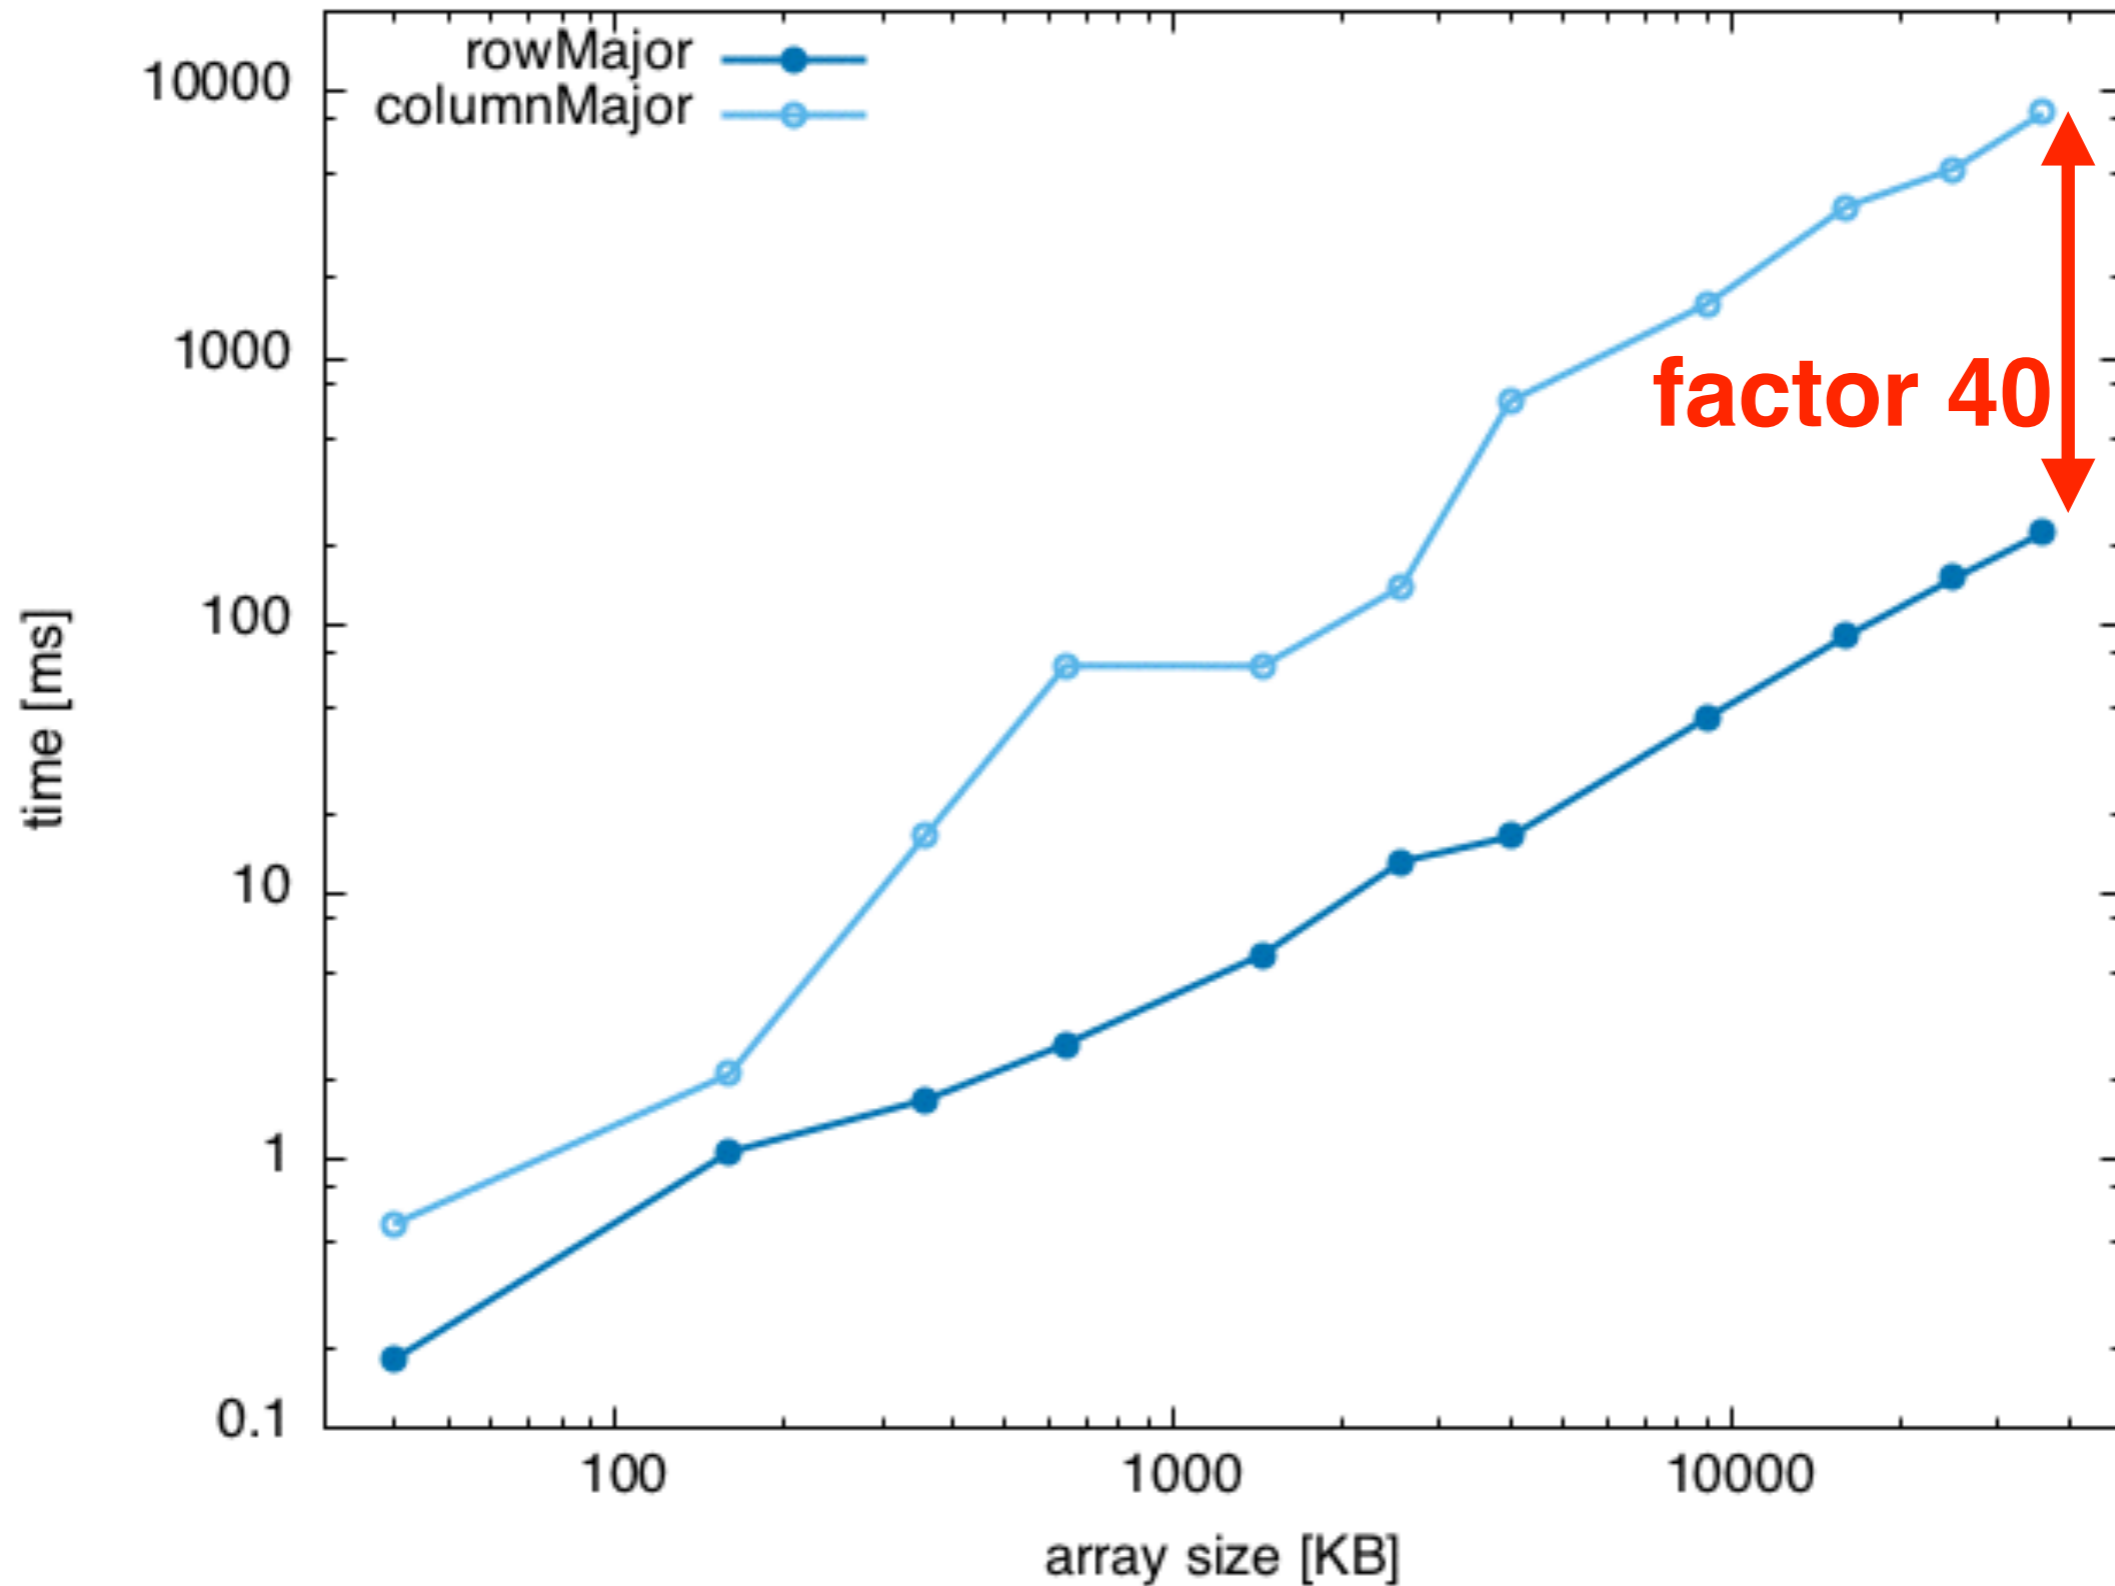
2D Array traversal

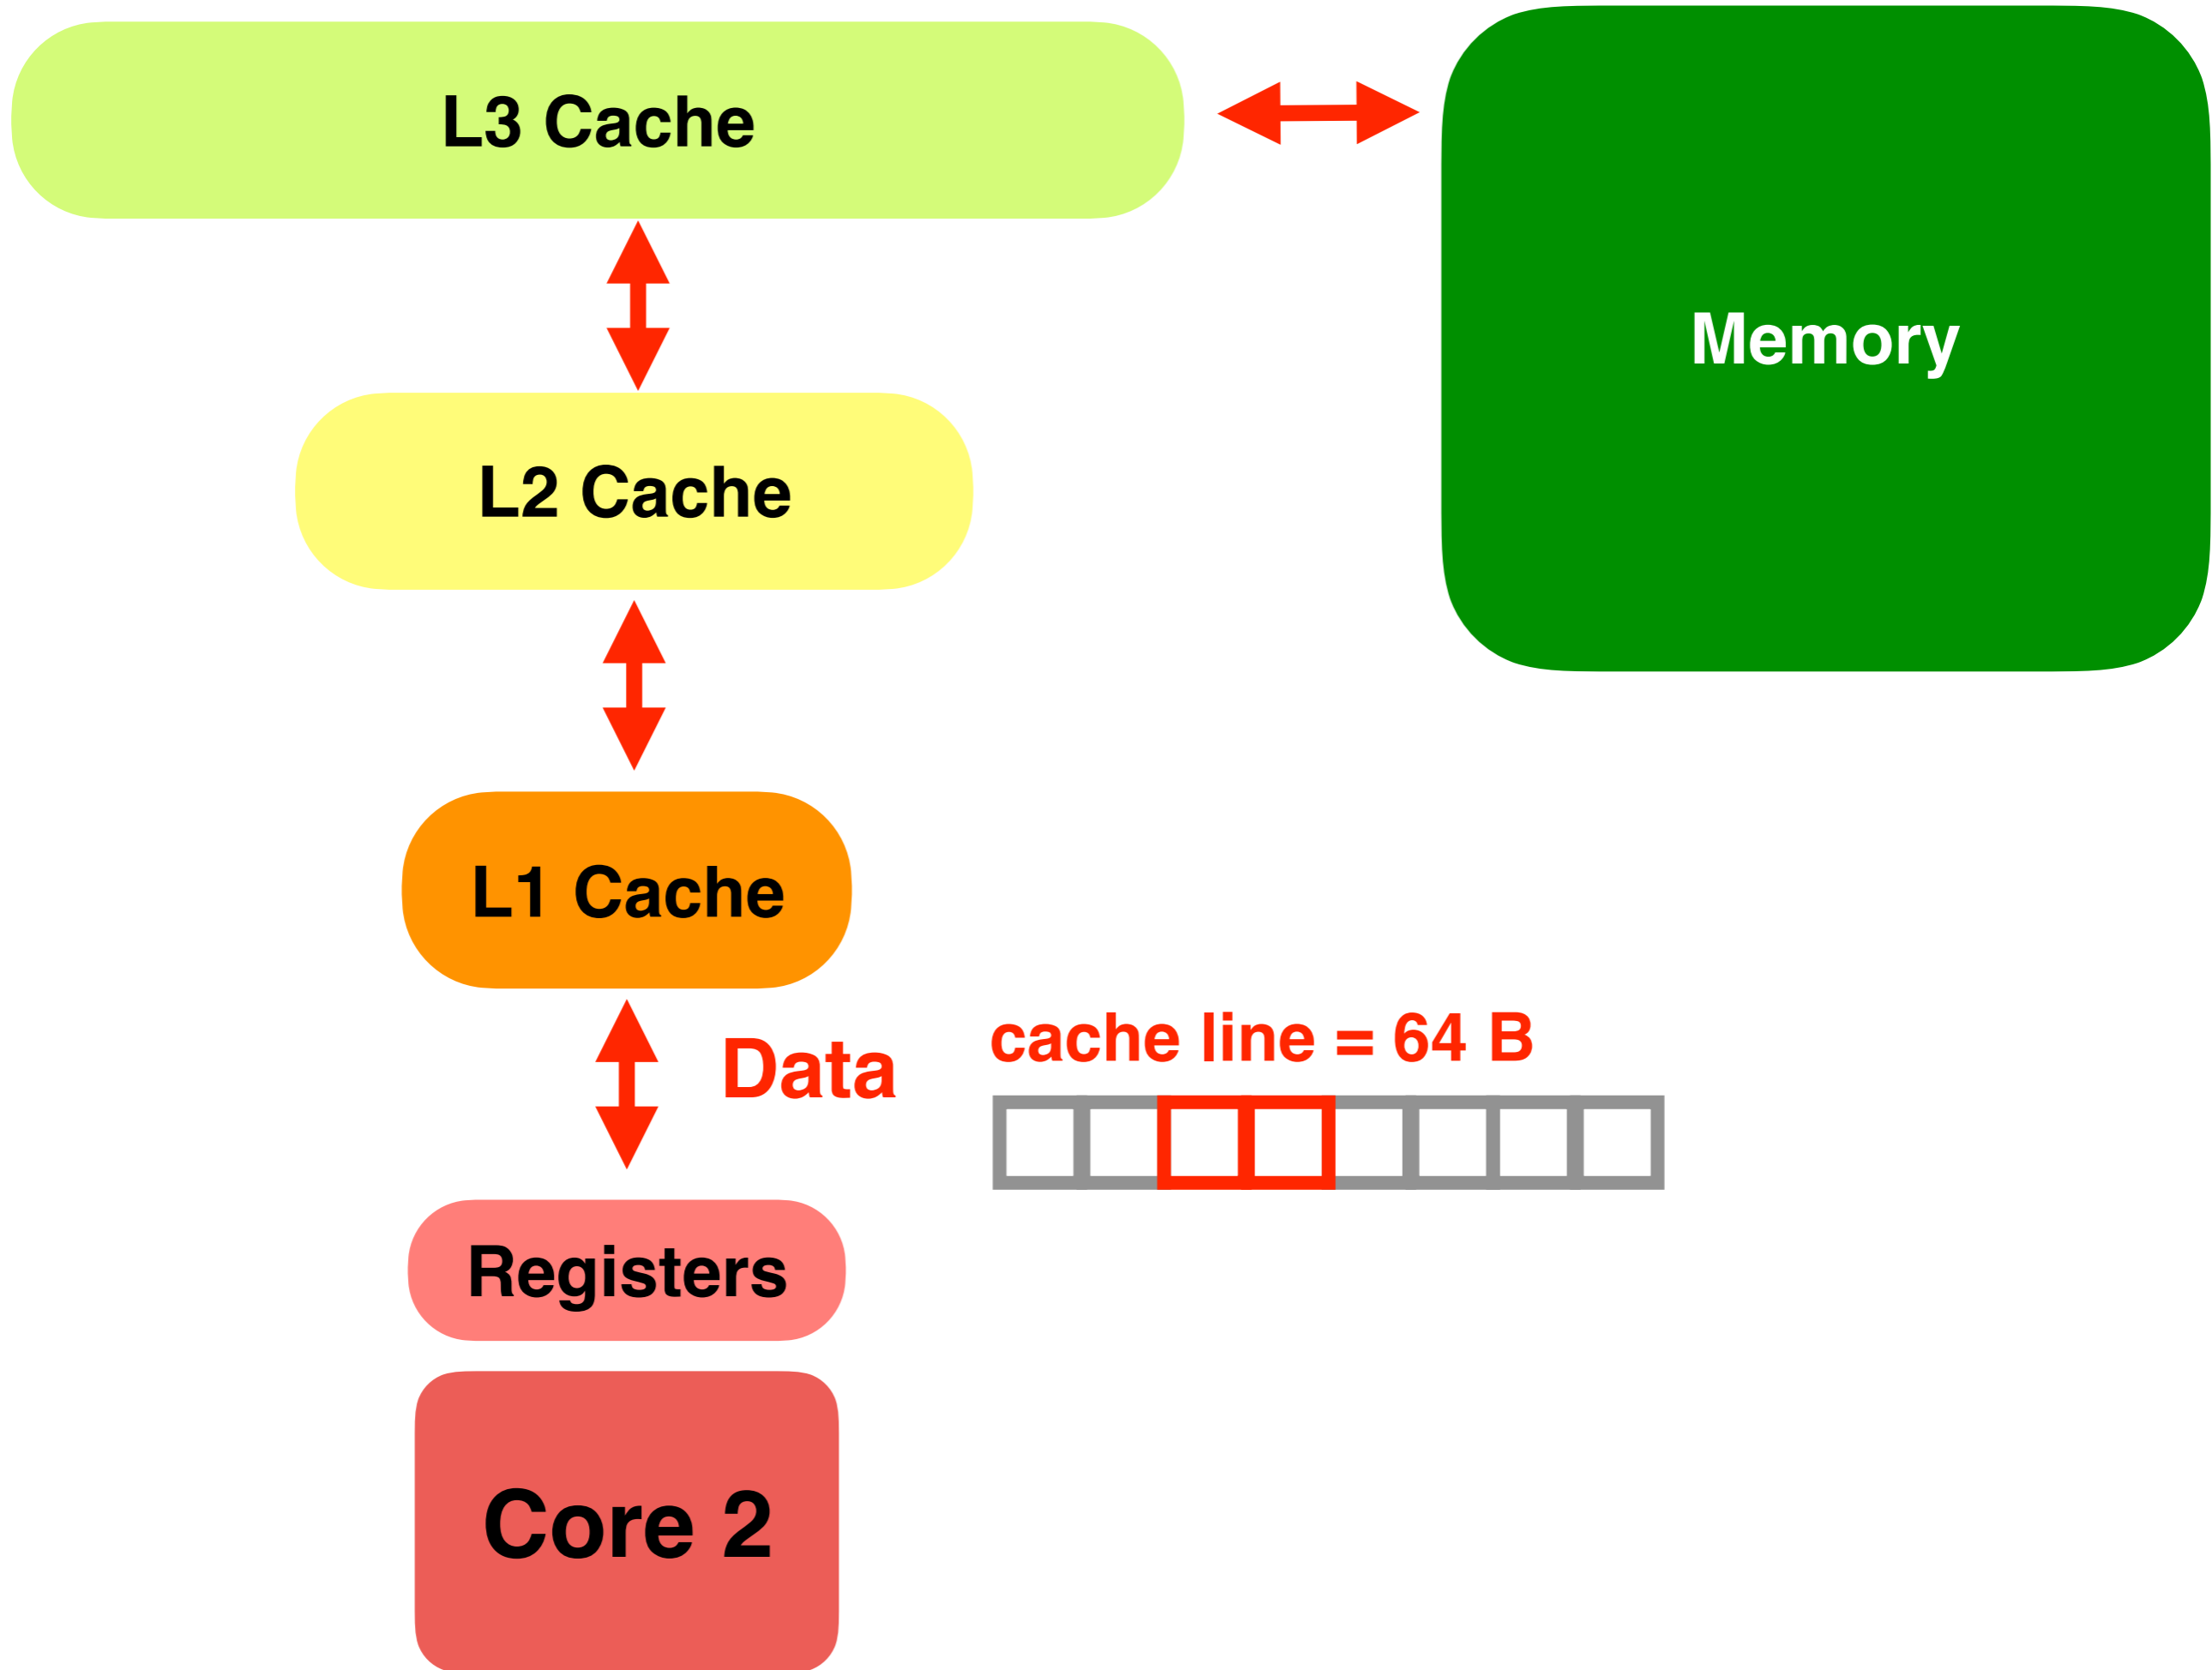
```
int array[n][n];
```

```
BENCHMARK_START(rowMajor, 100)
{
    for (int i = 0; i < n; ++i)
        for (int j = 0; j < n; ++j)
            array[i][j] += j;
}
BENCHMARK_STOP(rowMajor)
```

```
BENCHMARK_START(columnMajor, 100)
{
    for (int i = 0; i < n; ++i)
        for (int j = 0; j < n; ++j)
            array[j][i] += j;
}
BENCHMARK_STOP(columnMajor)
```

2D Array traversal benchmark





2D Array traversal + some work

```
float array[n][n];
```

```
BENCHMARK_START(rowMajor, 100)
```

```
{
```

```
    for (int i = 0; i < n; ++i)
```

```
        for (int j = 0; j < n; ++j)
```

```
            array[i][j] += std::sqrt (std::hash<int>() (j * n + i));
```

```
}
```

```
BENCHMARK_STOP(rowMajor)
```

```
BENCHMARK_START(columnMajor, 100)
```

```
{
```

```
    for (int i = 0; i < n; ++i)
```

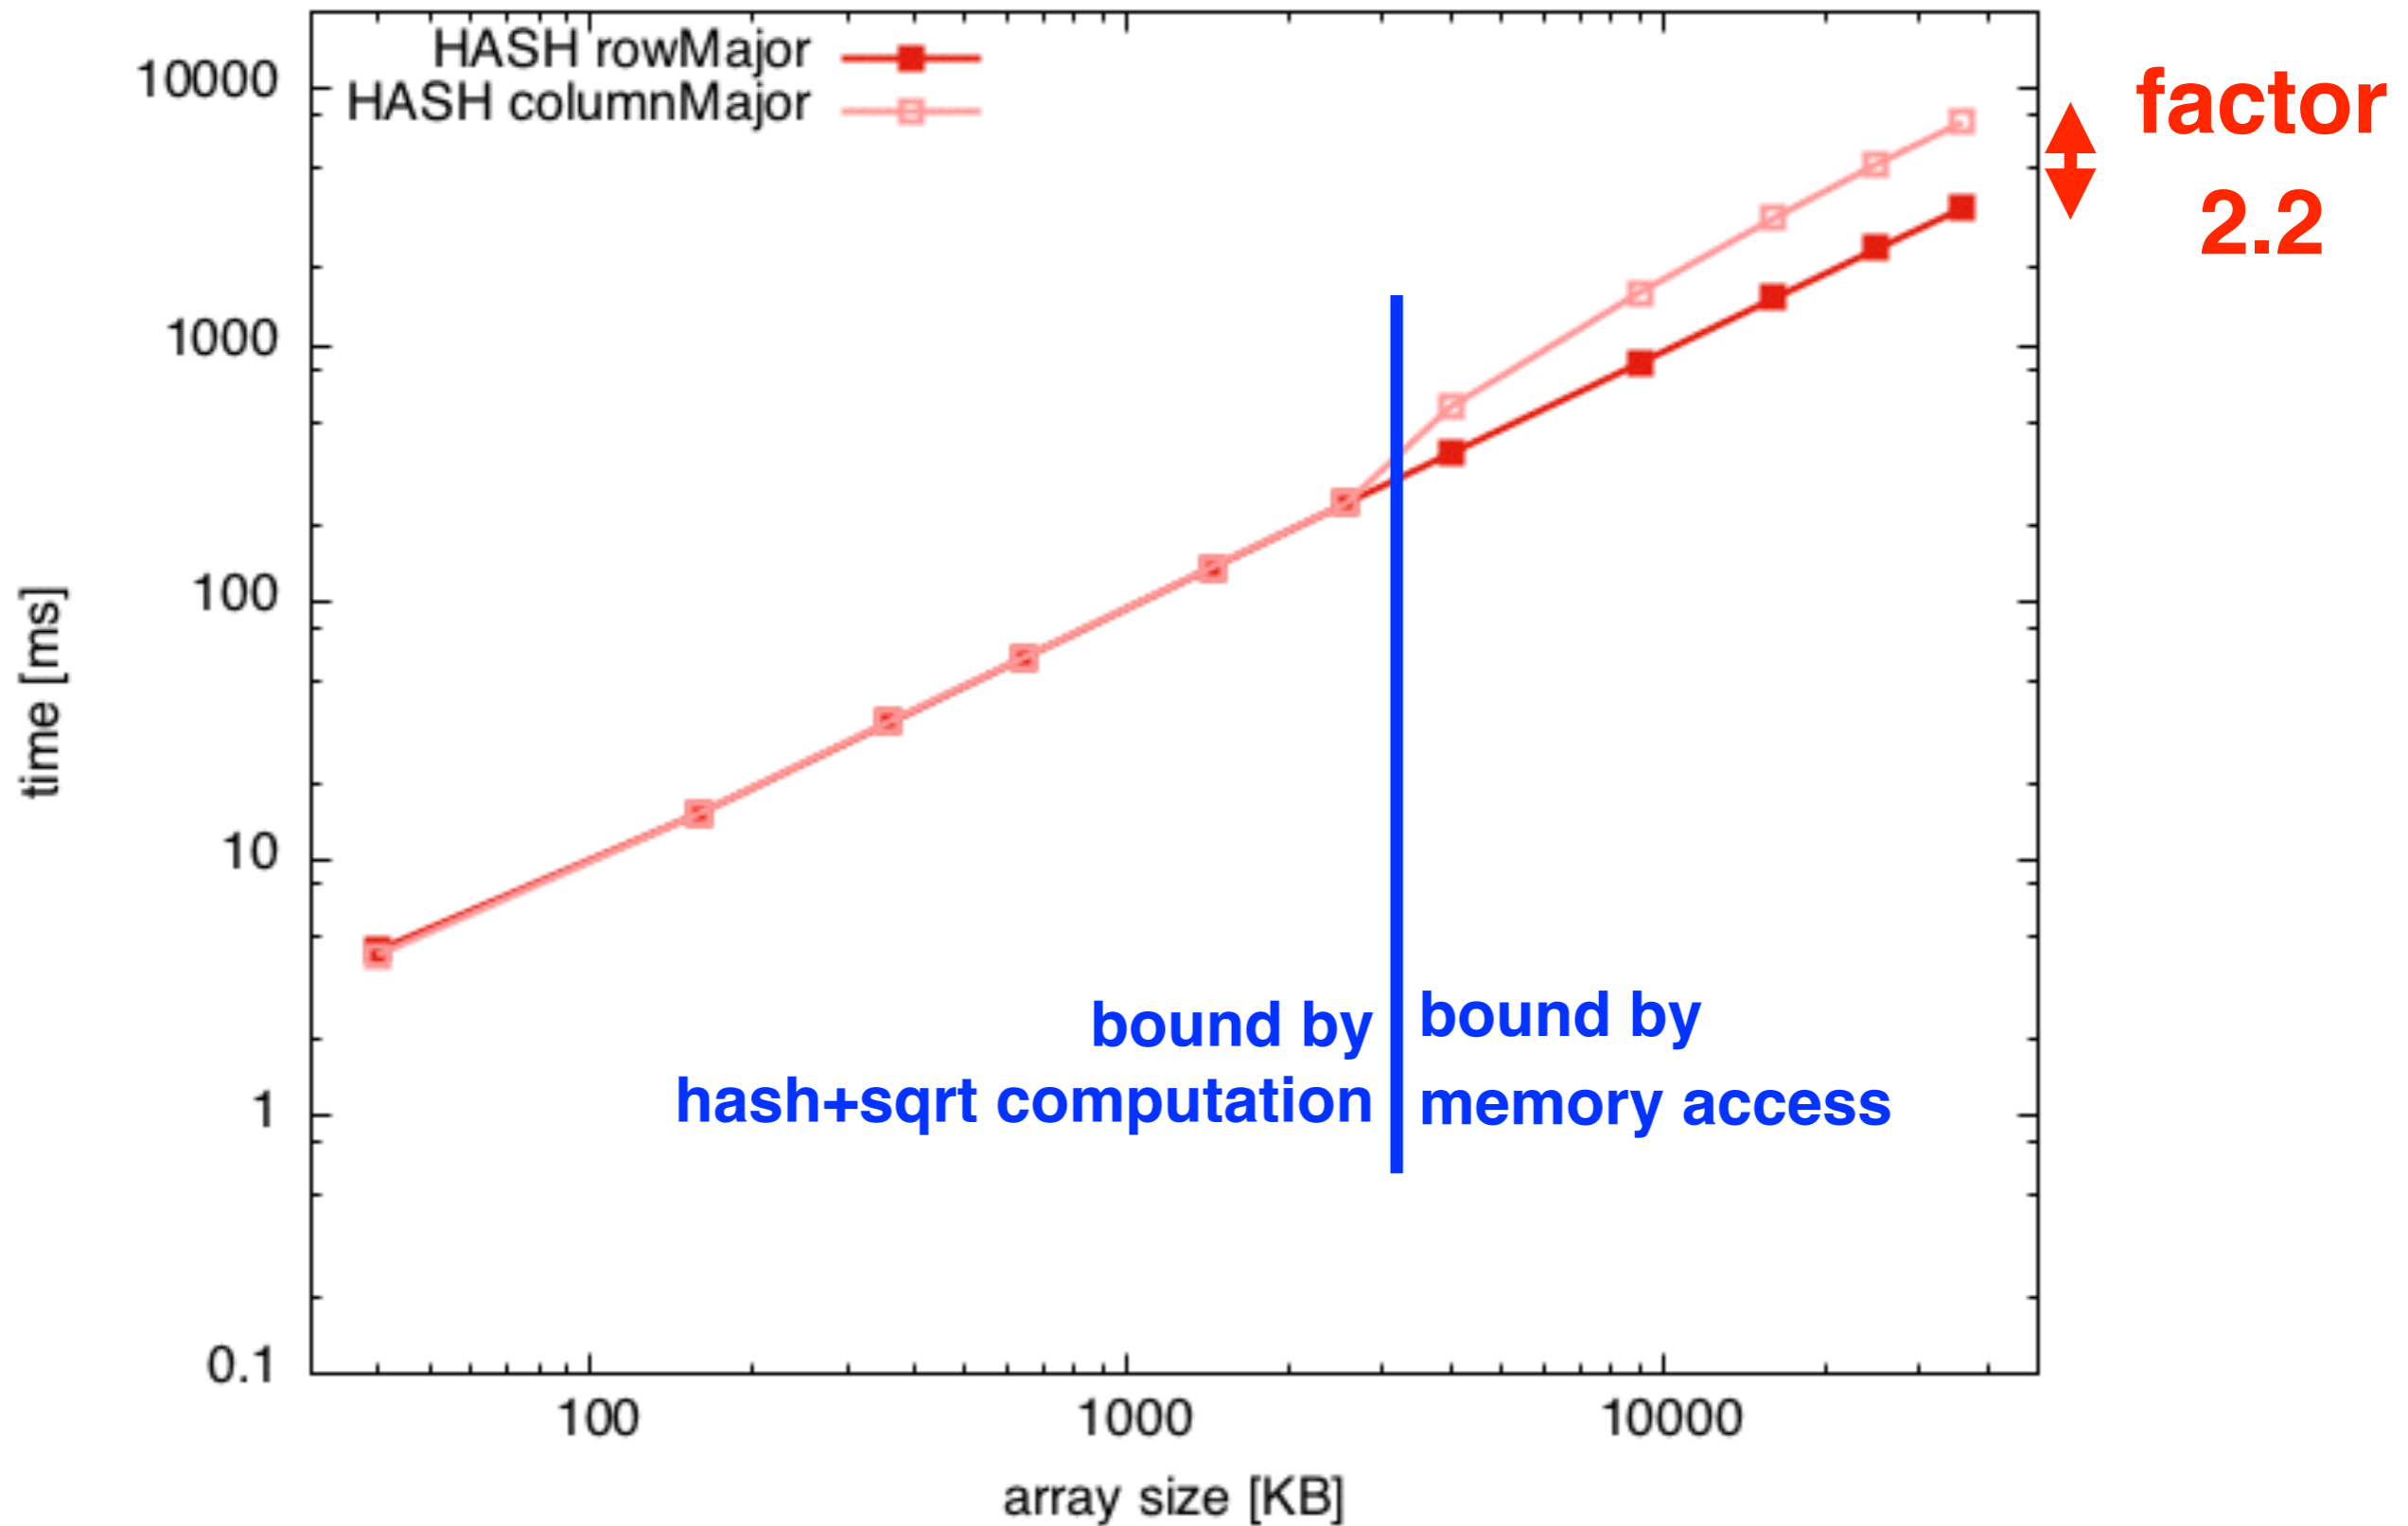
```
        for (int j = 0; j < n; ++j)
```

```
            array[j][i] += std::sqrt (std::hash<int>() (j * n + i));
```

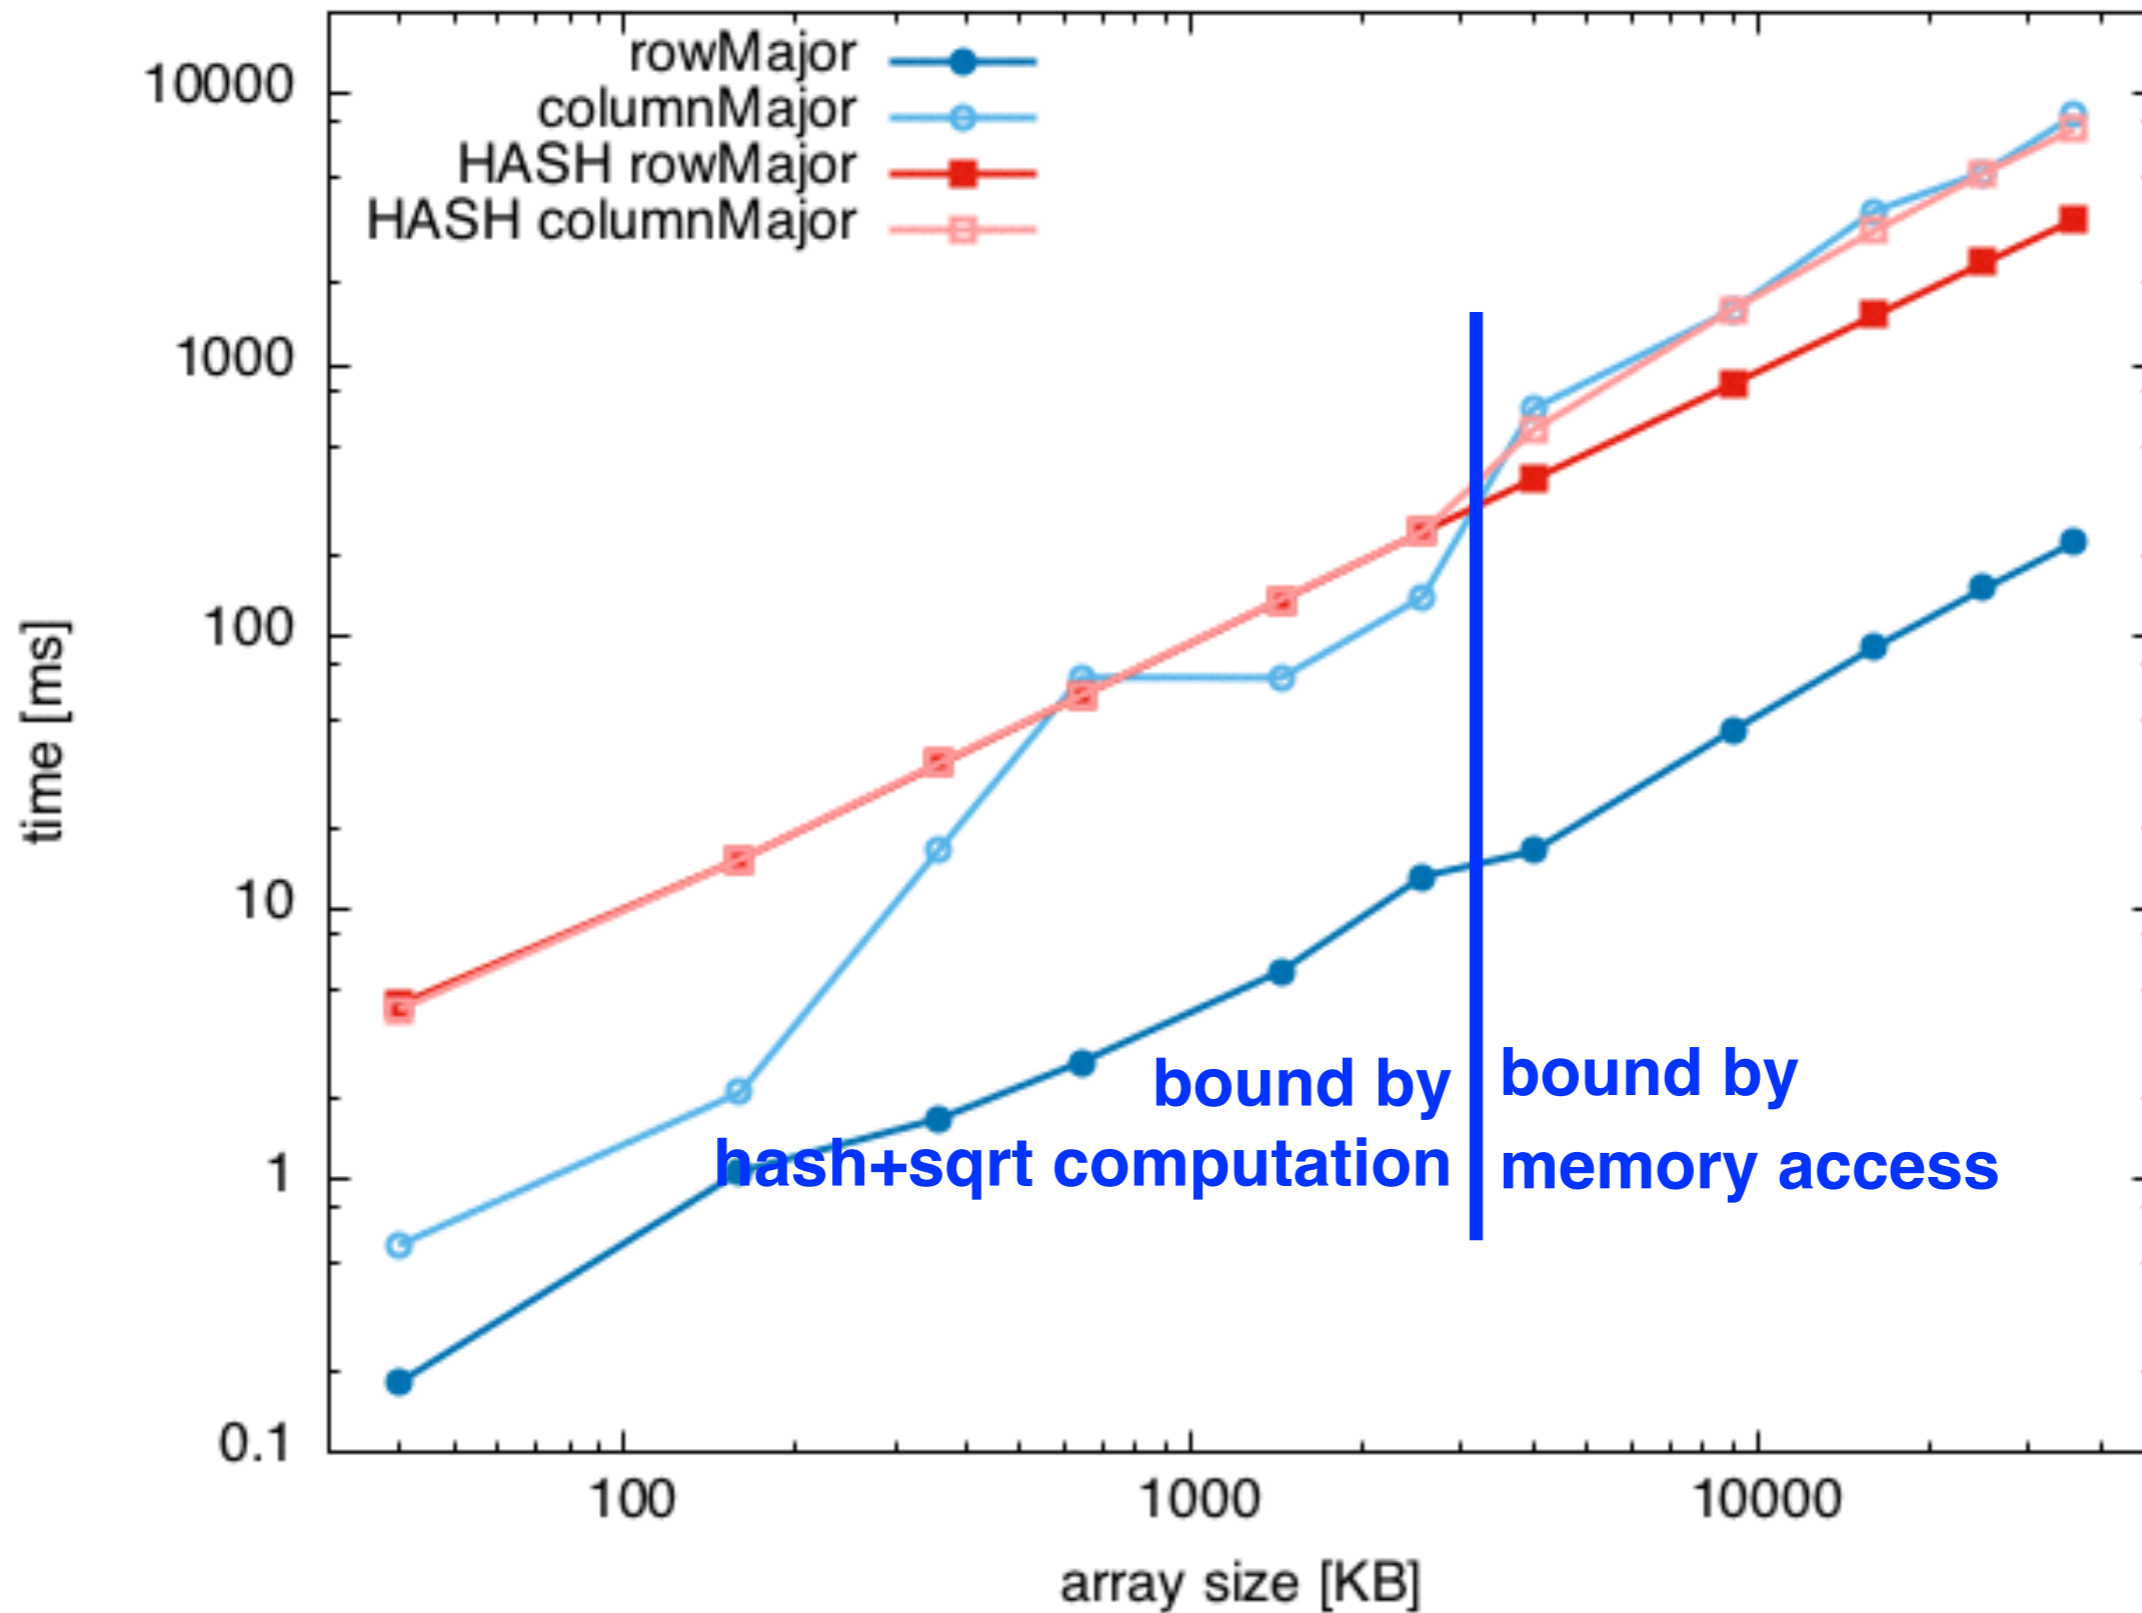
```
}
```

```
BENCHMARK_STOP(columnMajor)
```

2D Array traversal benchmark



2D Array traversal benchmark



2D Array traversal: time profile

(Xcode Instruments)

ails > Call Tree > Call Tree

| Running Time | Self (ms) | Symbol Name |
|-----------------|-----------|--|
| 4054.0ms 100.0% | 0.0 | ▼Main Thread 0x127ea |
| 4054.0ms 100.0% | 0.0 | ▼start libdyld.dylib |
| 4054.0ms 100.0% | 3489.0 | ▼main CacheProfileTest |
| 565.0ms 13.9% | 565.0 | std::__1::enable_if<is_integral<unsigned long>::value, double>::type std::__1::sqrt<un |

2D Array traversal + some work

```
float array[n][n];
```

```
BENCHMARK_START(rowMajor, 100)
```

```
{
```

```
    for (int i = 0; i < n; ++i)
```

```
        for (int j = 0; j < n; ++j)
```

```
            array[i][j] += std::sqrt (std::hash<int>() (j * n + i));
```

```
}
```

```
BENCHMARK_STOP(rowMajor)
```

```
BENCHMARK_START(columnMajor, 100)
```

```
{
```

```
    for (int i = 0; i < n; ++i)
```

```
        for (int j = 0; j < n; ++j)
```

```
            array[j][i] += std::sqrt (std::hash<int>() (j * n + i));
```

```
}
```


```
BENCHMARK_STOP(columnMajor)
```


2D Array traversal: time profile


(Xcode Instruments)


ails > Call Tree > Call Tree

| Running Time | Self (ms) | Symbol Name |
|-----------------|-----------|--|
| 4054.0ms 100.0% | 0.0 | ▼Main Thread 0x127ea |
| 4054.0ms 100.0% | 0.0 | ▼start libdyld.dylib |
| 4054.0ms 100.0% | 3489.0 | ▼main CacheProfileTest |
| 565.0ms 13.9% | 565.0 | std::__1::enable_if<is_integral<unsigned long>::value, double>::type std::__1::sqrt<un |










time



time

```
struct Foo
{
    char c;
    double d;
    short s;
    int i;
};
```

```
class Bar
{
    uint64_t lo;
    uint64_t mid;
    uint64_t hi;
};
```

```
std::vector<Foo> foos (1000);
std::vector<Bar> bars (1000);
```


```
doSomething (foos[i], bars[i]);
```

```
class Baz
{
    Foo foo;
    Bar bar;
};
```

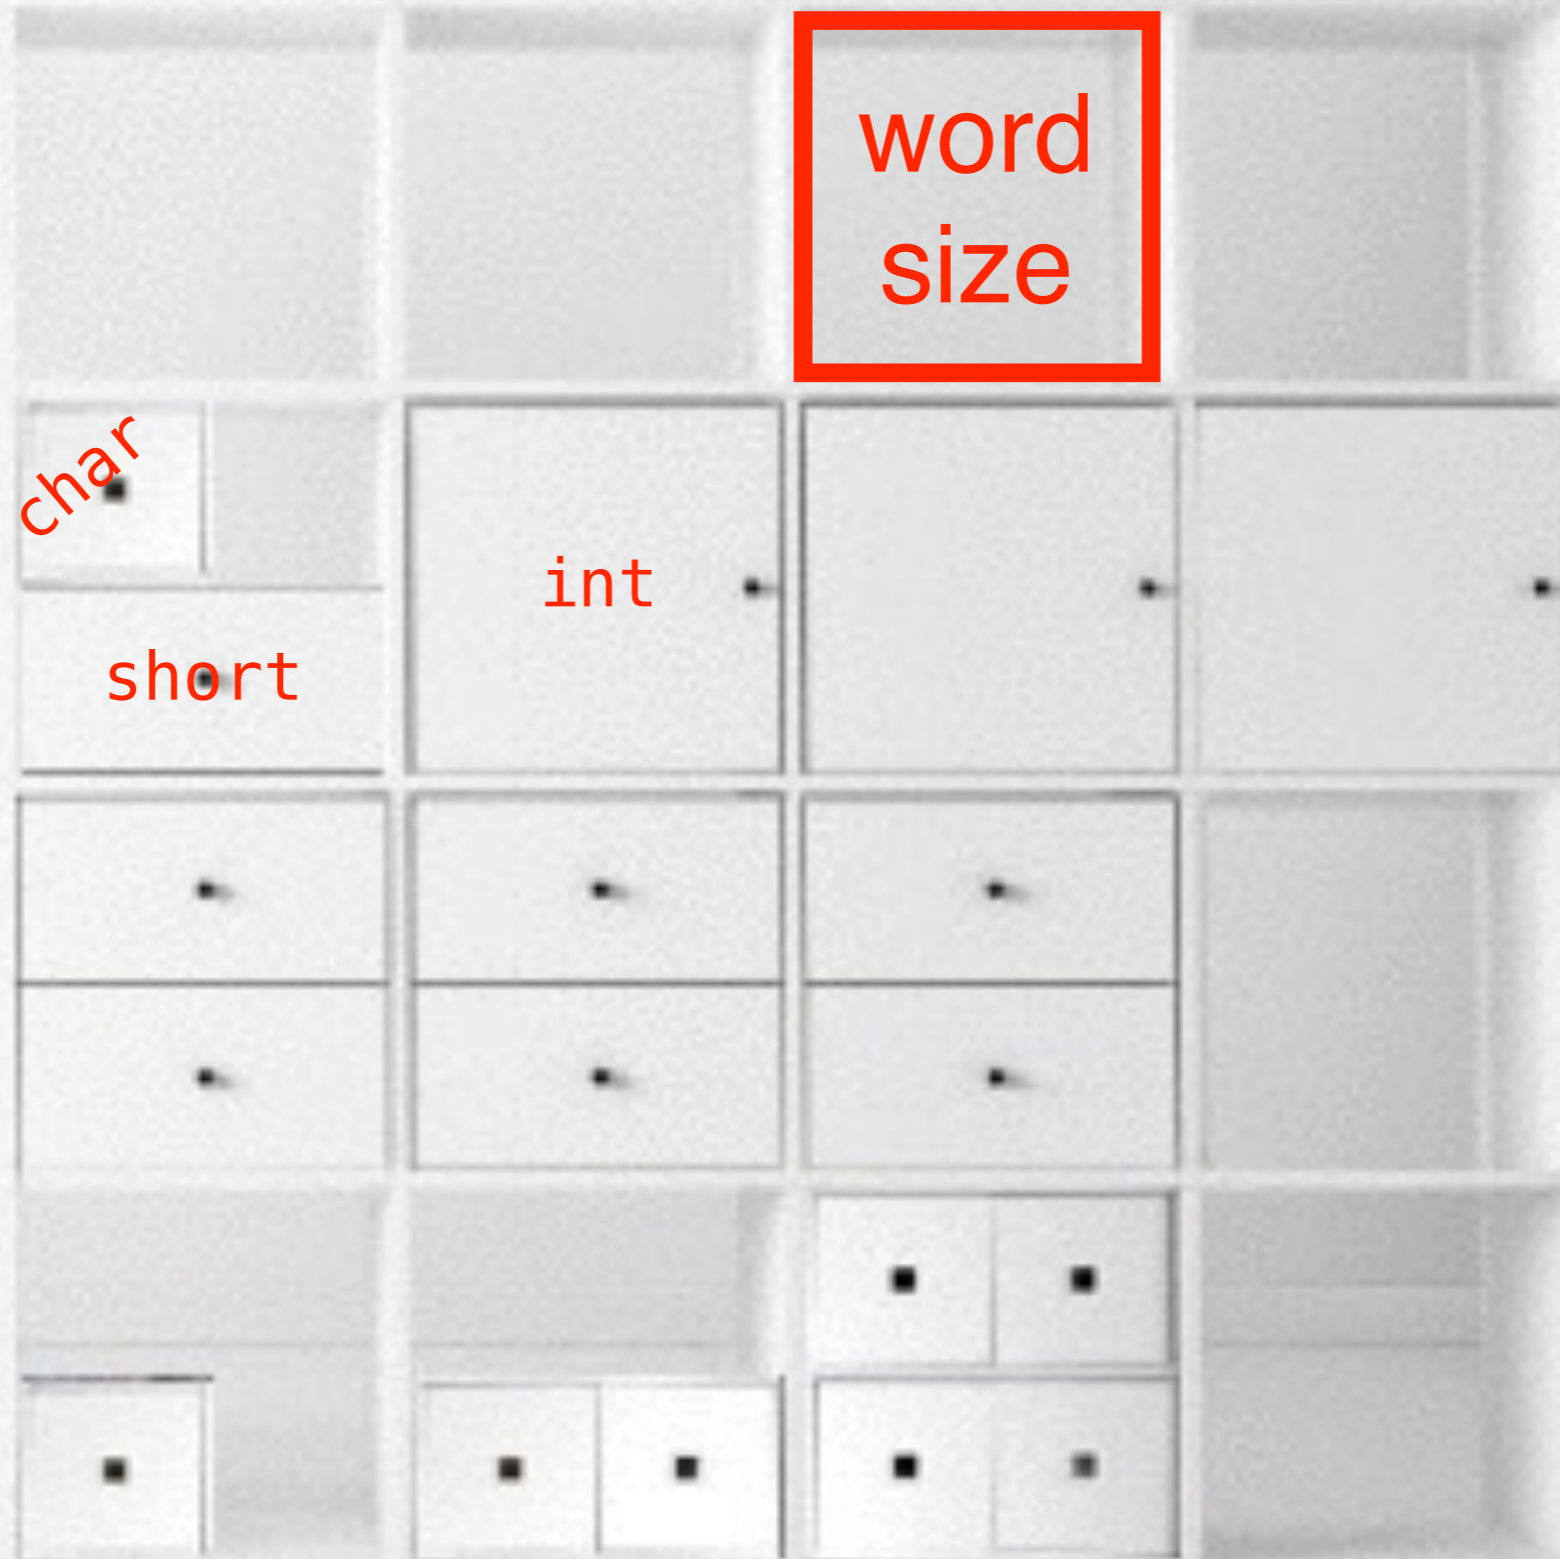
```
std::vector<Baz> bazz (1000);
```

```
doSomething (bazz[i]);
```

memory
alignment



ALIGNOF




Alignment requirements

(on mainstream desktop & mobile platforms)


| <code>T</code> | <code>sizeof (T)</code> | <code>alignof (T)</code> |
|-------------------------------|-------------------------|--------------------------|
| <code>char</code> | 1 | 1 |
| <code>short</code> | 2 | 2 |
| <code>int</code> | 4 | 4 |
| <code>uint64_t</code> | 8 | 8 |
| <code>long double</code> | 8 or 16 | 8 or 16 |
| <code>std::max_align_t</code> | 8 or 16 | 8 or 16 |

Alignment of user-defined classes

```
struct Foo
{
    char c;
    double d;
    short s;
    int i;
};
```




```
struct Foo
{
    char c;
    short s;
    int i;
    double d;
};
```




Packed structs

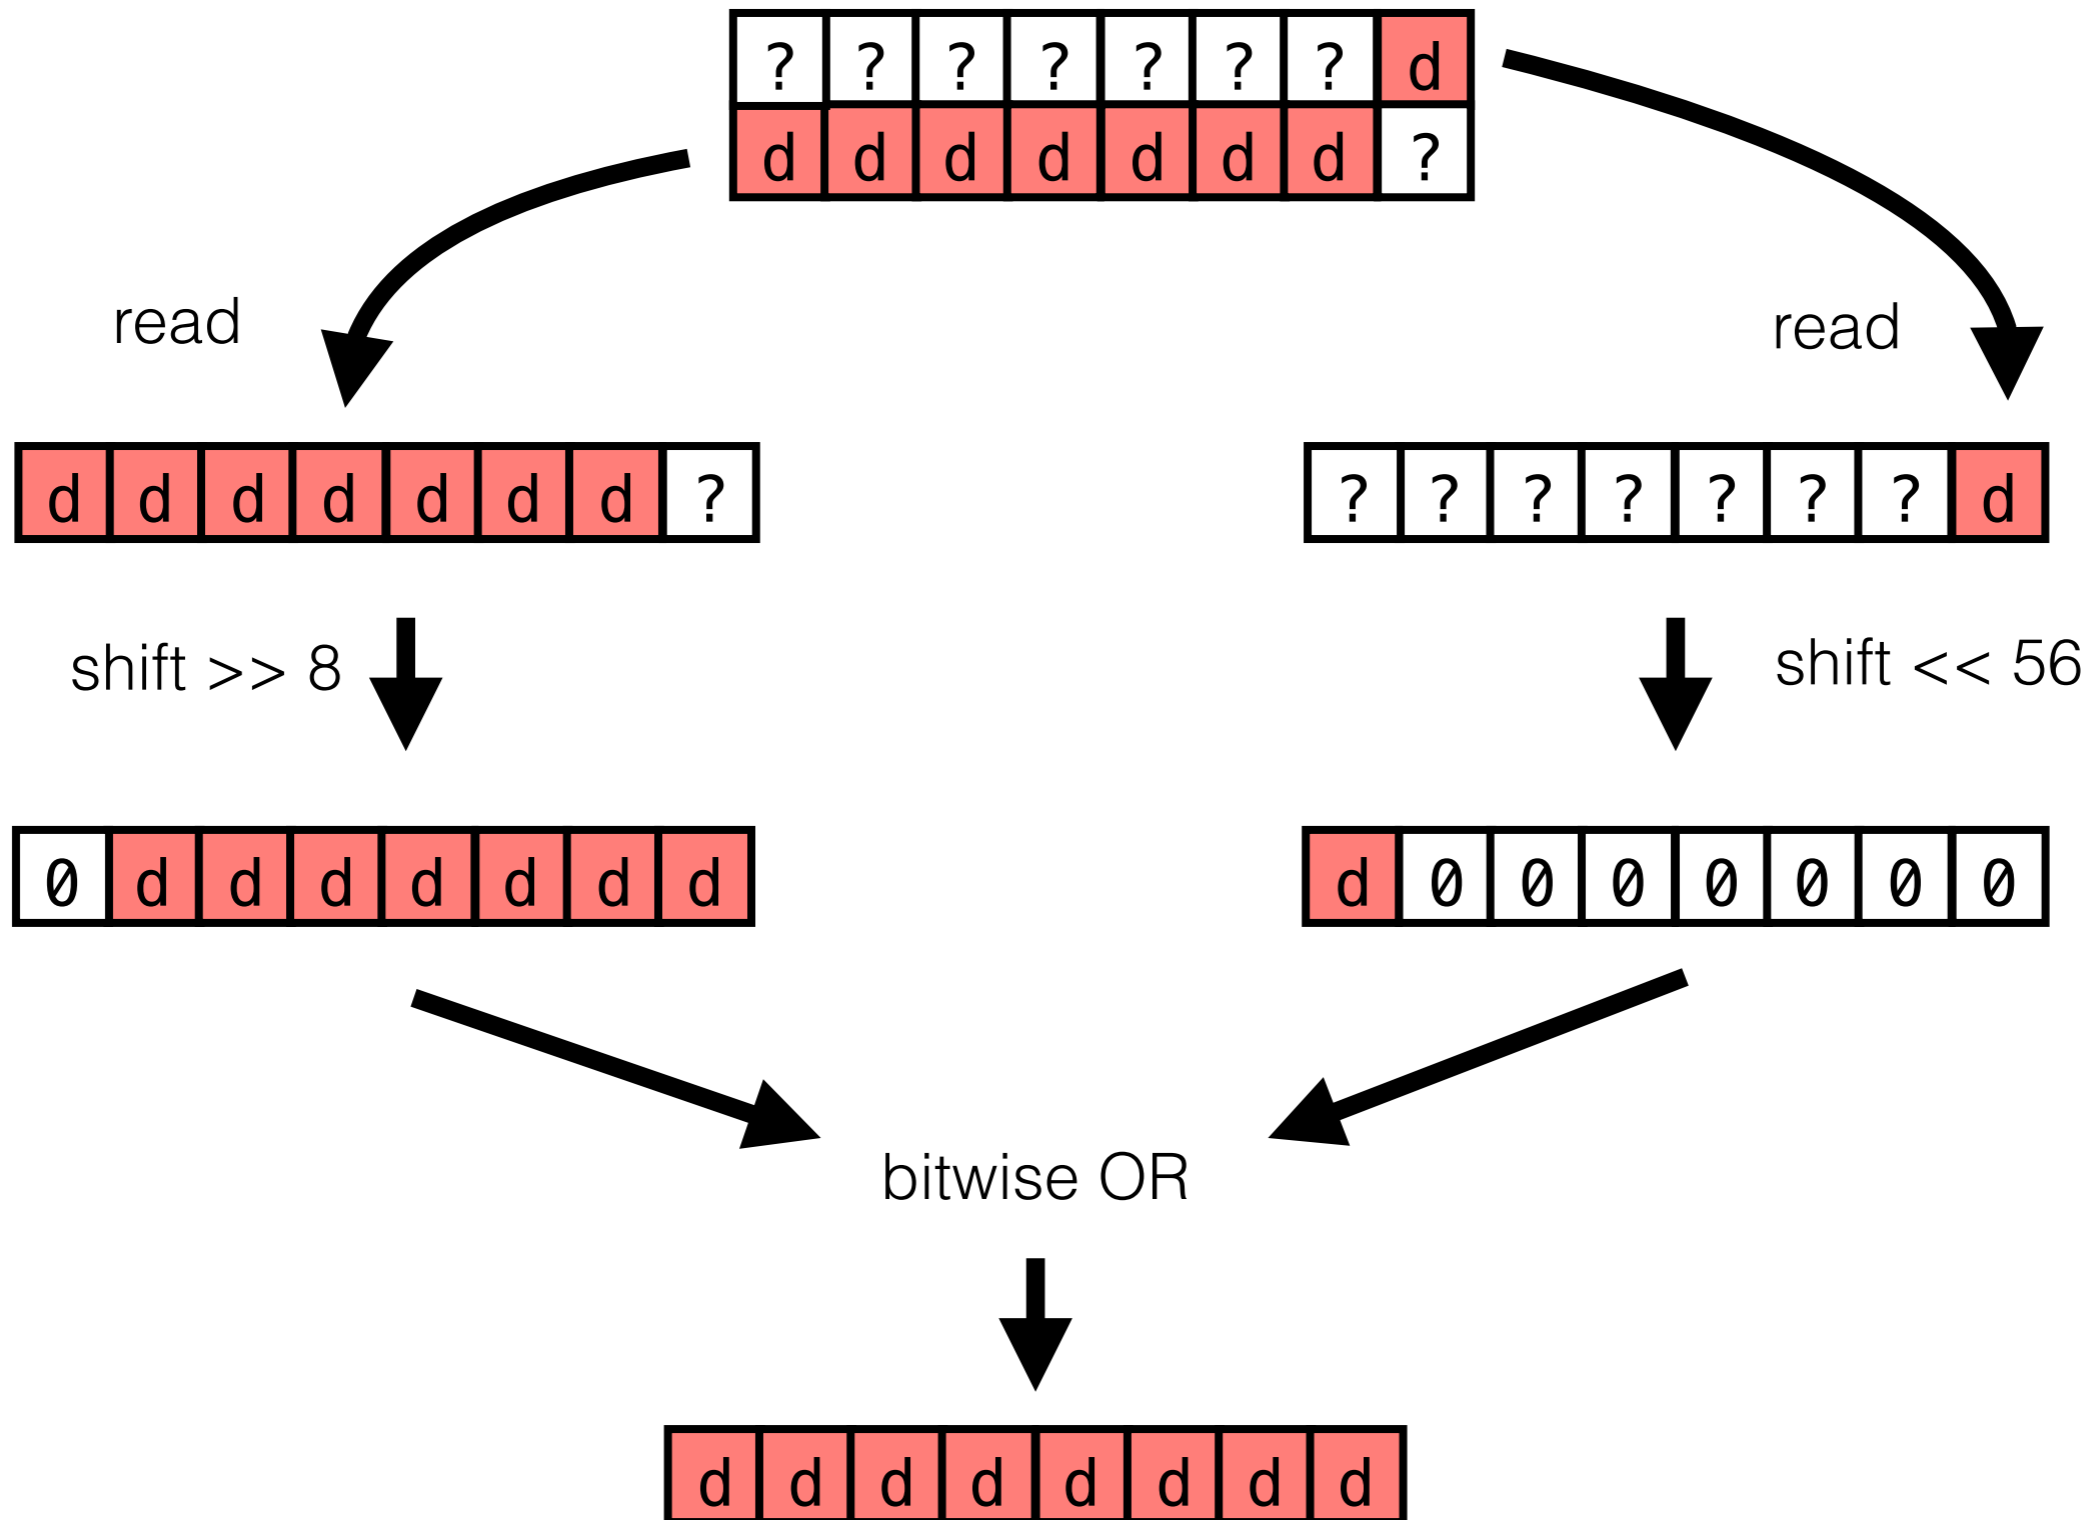
```
struct Foo
{
    char c;
    short s;
    int i;
    double d;
}
__attribute__((packed));
```



```
struct BWAChunk
{
    char    description[256];
    char    originator[32];
    char    originatorRef[32];
    char    originationDate[10];
    char    originationTime[8];
    uint32_t timeRefLow;
    uint32_t timeRefHigh;
    uint16_t version;
    uint8_t  umid[64];
    uint8_t  reserved[190];
    char     codingHistory[1];
}
__attribute__((packed));
```



Unaligned memory access



Unaligned access benchmark



i7



Core2Duo



iPhone6S

AppleA9



Nexus5X

Snapdragon 808


```
clang -x c++ -std=c++11 -stdlib=libc++ -O3
```

Unaligned access benchmark


```
template <typename T>
struct AlignedStruct
{
    char c;
    T value;
};
```

```
template <typename T>
struct PackedStruct
{
    char c;
    T value;
}
__attribute__((packed));
```

AlignedStruct<double>



PackedStruct<double>



Unaligned access benchmark

```
template <typename StructT>
auto test (std::size_t size)
{
    std::vector<float> vecFloats (size);

    std::iota (vecFloats.begin(),
               vecFloats.end(),
               0.0);

    std::vector<StructT> vecStructs (size);

    std::generate (vecStructs.begin(),
                   vecStructs.end(),
                   [] { return StructT {
                       'x', double (rand() % 100)};
                   });

    BENCHMARK_START(loop, 10000)

    for (int i = 0; i < size; ++i)
        vecStructs[i].value += vecFloats[i];

    BENCHMARK_STOP(loop)

    return vecStructs;
}
```

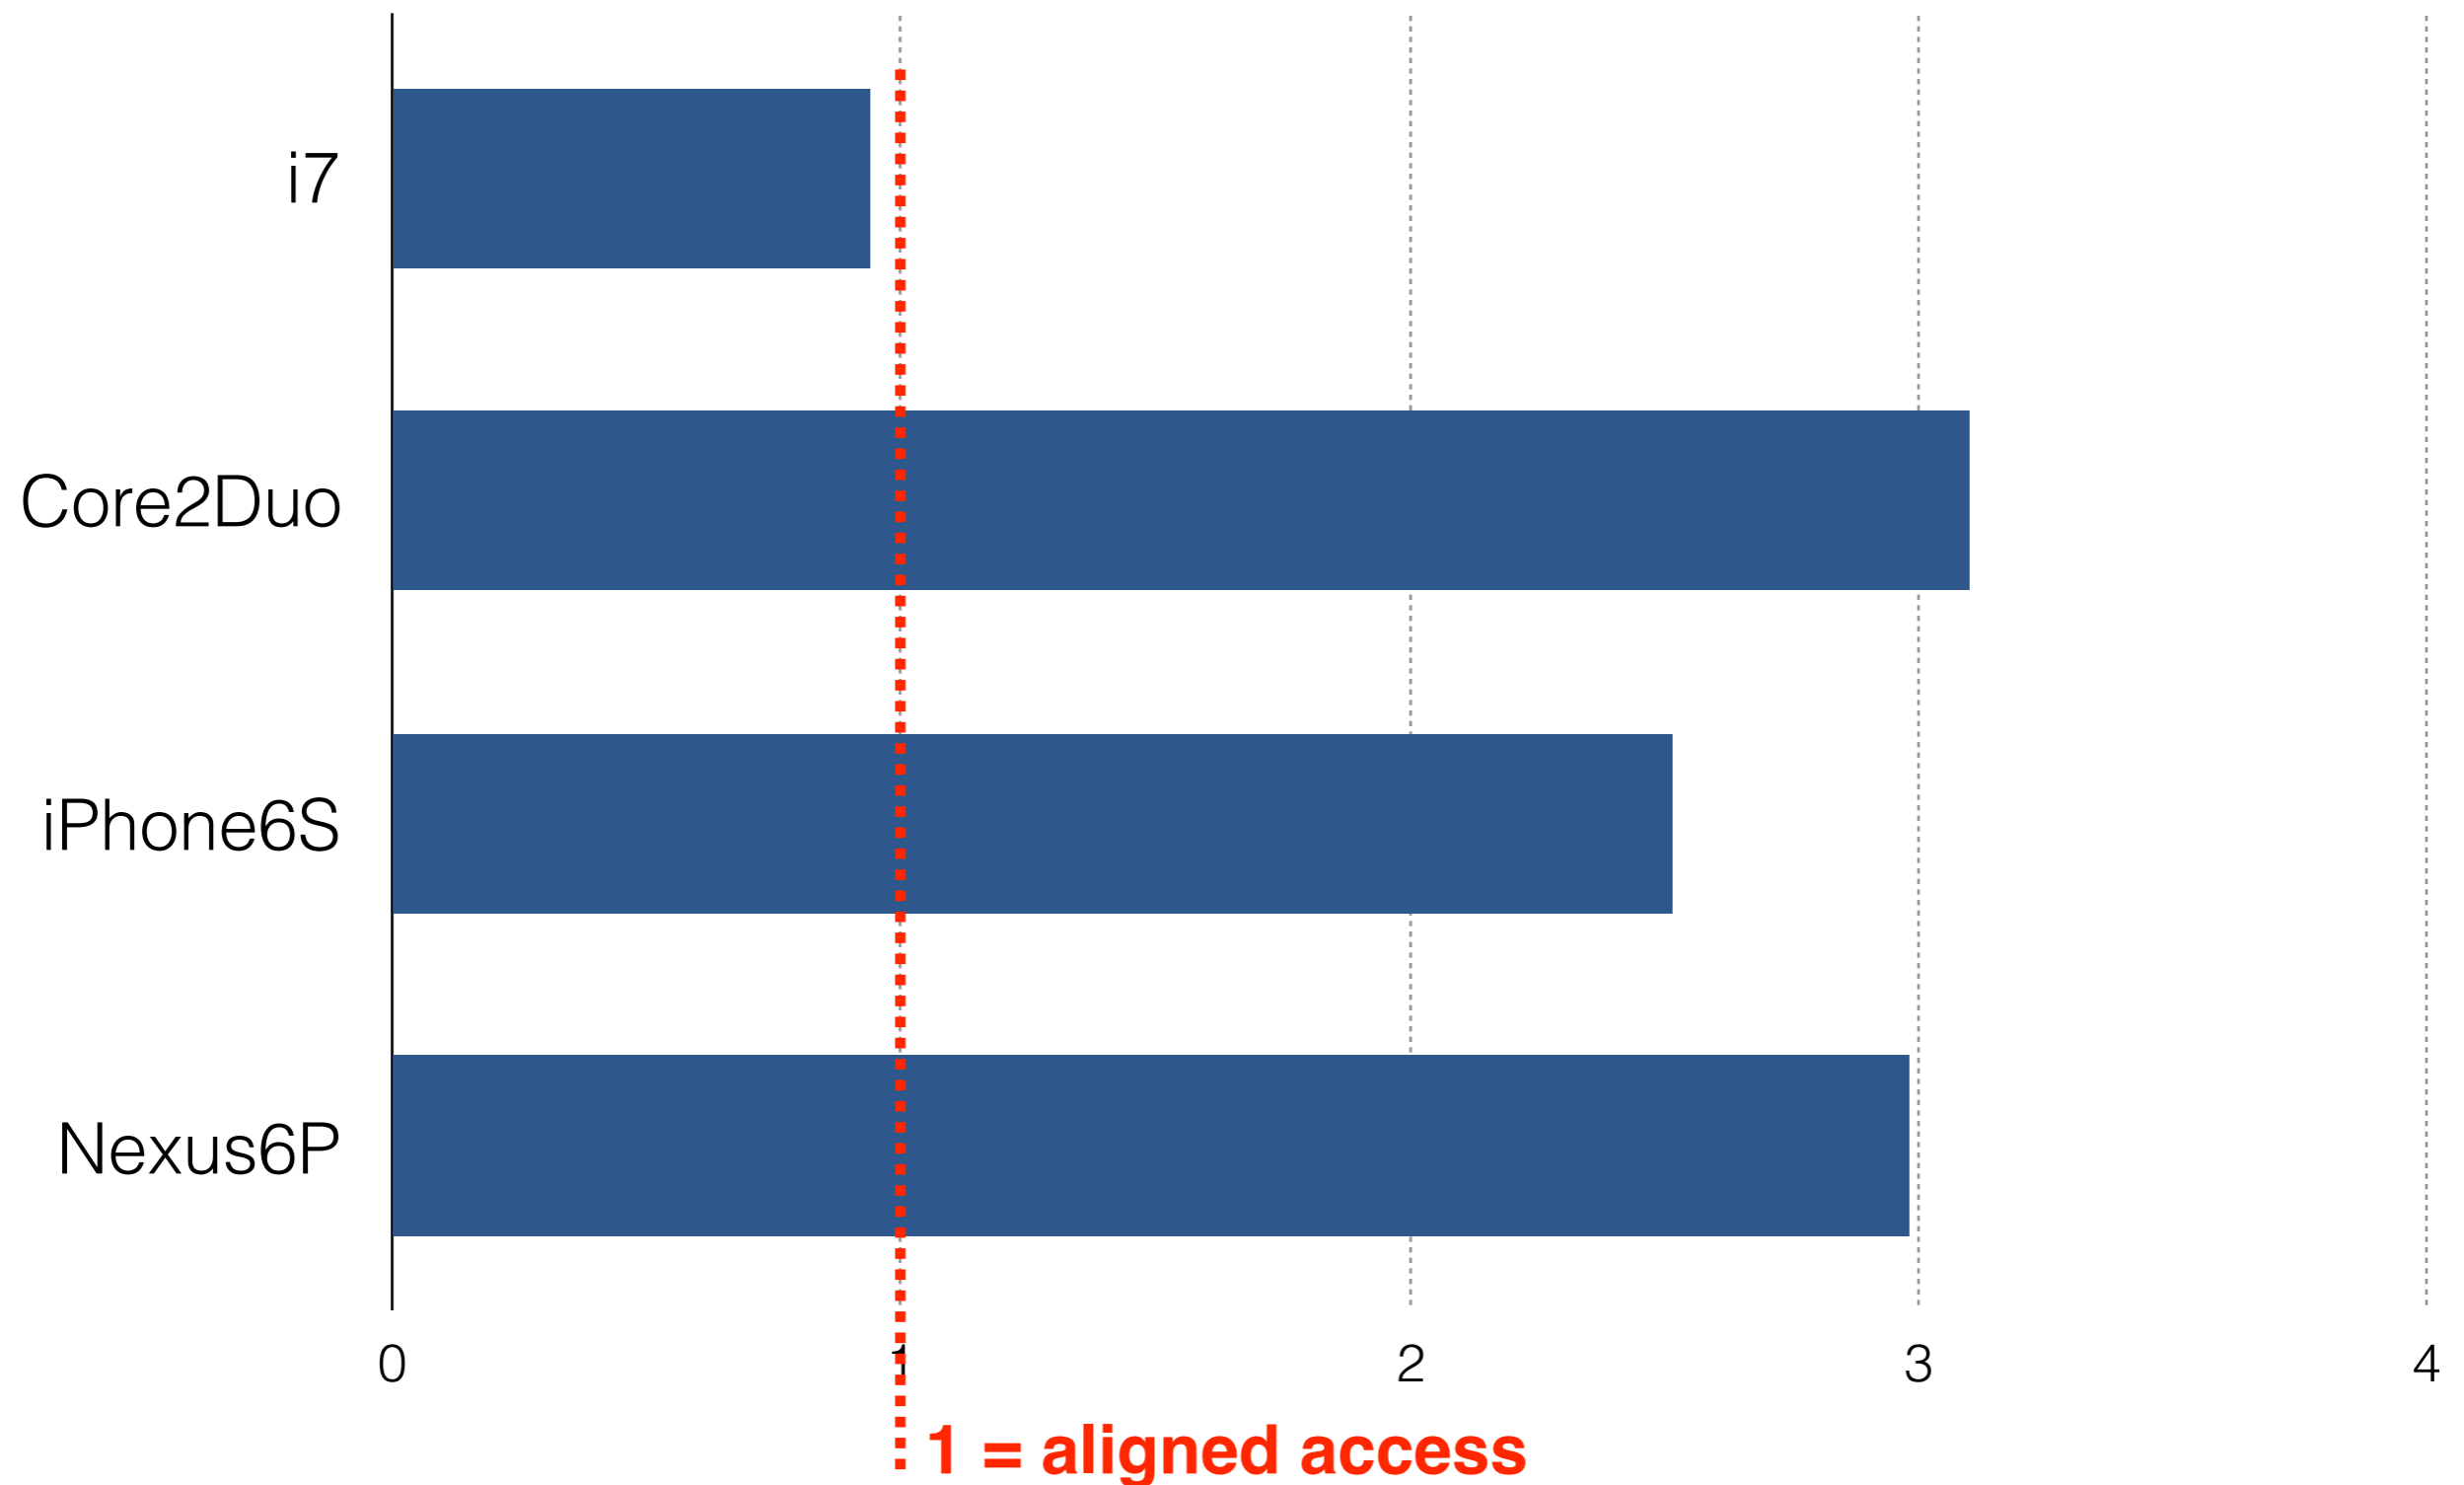
AlignedStruct<double>

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| c | | | | | | | |
| d | d | d | d | d | d | d | d |
| c | | | | | | | |
| d | d | d | d | d | d | d | d |
| c | | | | | | | |
| d | d | d | d | d | d | d | d |

PackedStruct<double>

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| c | d | d | d | d | d | d | d |
| d | c | d | d | d | d | d | d |
| d | d | c | d | d | d | d | d |
| d | d | d | c | d | d | d | d |
| d | d | d | d | c | d | d | d |
| d | d | d | d | d | c | d | d |


Unaligned access benchmark (double offset by 1 byte)



Casts between differently-aligned types

```
struct Foo // alignof (Foo) == 1
{
    char c[8];

    double interpretAsDouble() noexcept
    {
        return *(reinterpret_cast<double*> (c));
    }
};
```

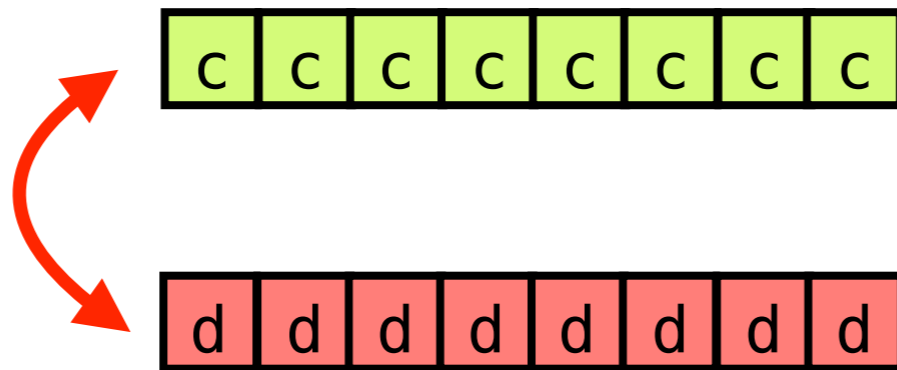


----- beginning of crash

Fatal signal 7 (SIGBUS), code 1, fault addr 0xc20b8009

Casts between differently-aligned types

```
union Foo    // alignof (Foo) == 8
{
    char c[8];
    double d;
};
```



reinterpret_cast vs. union

```
uint64_t doubleAsInt (double x) noexcept // version 1
{
    union { double asDouble; uint64_t asInt; } u;
    u.asDouble = x;
    return u.asInt;
}

uint64_t doubleAsInt (double x) noexcept // version 2
{
    return *(reinterpret_cast<uint64_t*> (&x));
}

int main()
{
    std::vector<double> vec(4000);
    std::generate (vec.begin(), vec.end(), [] { return double (rand() % 1000); });

    uint64_t sum;

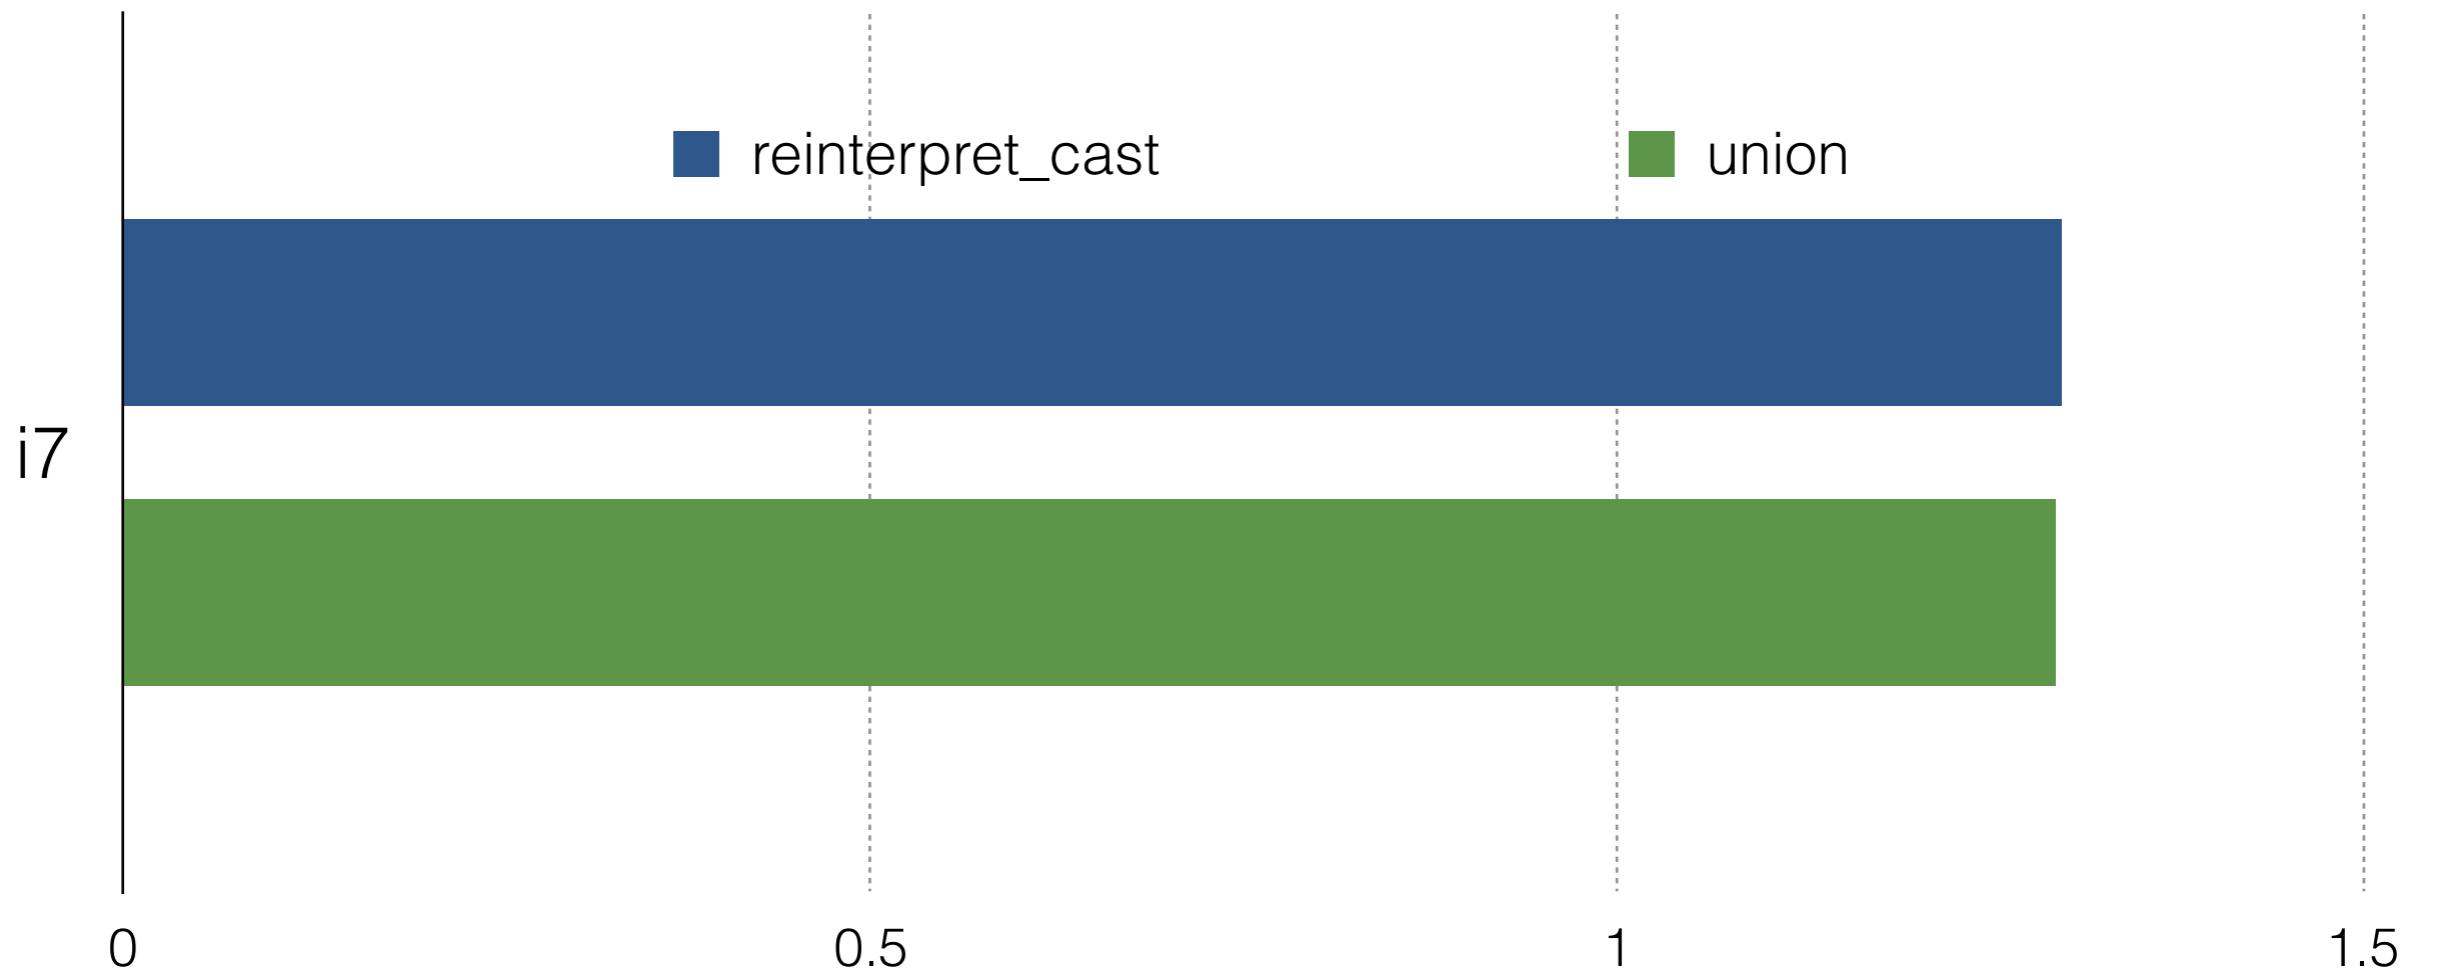
    BENCHMARK_START (doubleAsInt, 10000)

    for (auto& x : vec)
        sum += doubleAsInt (x);

    BENCHMARK_STOP (doubleAsInt)

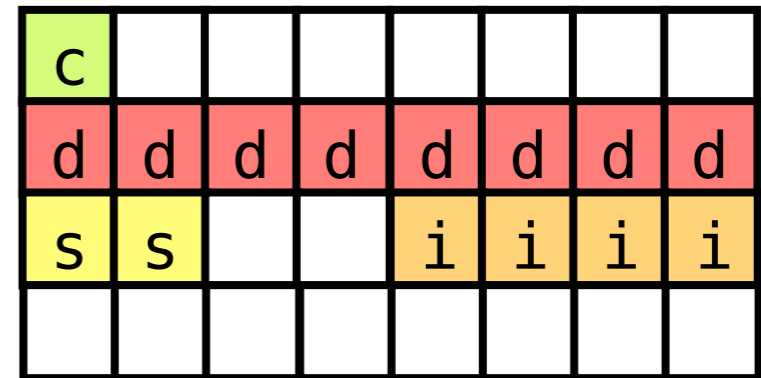
    std::cout << sum << '\n';
}
```

reinterpret_cast vs. union



Overalignment

```
struct alignas (32) Foo  
{  
    char c;  
    double d;  
    short s;  
    int i;  
};
```



Alignment on the heap

```
size_t bufferSize = 1024;  
void* scratchBuffer = new char[bufferSize];  
// ...and make this aligned by 8/16/32
```


Alignment on the heap

std::max_align_t

Defined in header `<cstddef>`

```
typedef /*implementation-defined*/ max_align_t; (since C++11)
```

std::max_align_t is a POD type whose [alignment requirement](#) is at least as strict (as large) as that of every scalar type.

Notes


Pointers returned by allocation functions such as `std::malloc` are suitably aligned for any object, which means they are aligned at least as strict as **std::max_align_t**.

std::max_align_t is usually synonymous with the largest scalar type, which is `long double` on most platforms, and its alignment requirement is either 8 or 16.

Alignment on the heap

```
auto* a = new char;  
auto* b = new char;  
auto* c = new char;
```


i7



Alignment on the heap

```
auto* a = new char;  
auto* b = new char;  
auto* c = new char;
```


Core2Duo



Alignment on the heap

```
auto* a = new char;  
auto* b = new char;  
auto* c = new char;
```

iPhone6S



Alignment on the heap

```
auto* a = new char;  
auto* b = new char;  
auto* c = new char;
```

Nexus5X



Alignment on the heap

```
struct alignas (8) Foo  
{  
    char ch;  
}
```

```
auto* a = new Foo;  
auto* b = new Foo;  
auto* c = new Foo;
```

Nexus5X



Alignment on the heap

```
template <size_t Size, size_t Align = Size>  
struct aligned_storage;
```

```
template <size_t Size, typename... Types>  
struct aligned_union;
```


Alignment on the heap

```
size_t bufferSize = 1024;
```

```
void* scratchBuffer = static_cast<void*> (  
    new std::aligned_storage<32, 32>::type[1024 / 32]);
```



Alignment on the heap

```
size_t bufferSize = 1024;
size_t bufferSizeWithPadding = bufferSize + 32;

void* scratchBuffer = new char[bufferSizeWithPadding];

if (! std::align (32, bufferSize,
                 scratchBuffer, bufferSizeWithPadding))
{
    // handle...
}
```

posix_memalign(3) - Linux man page

Name

posix_memalign, aligned_alloc, memalign, valloc, pvalloc - allocate aligned memory

Synopsis

```
#include <stdlib.h>
```

```
int posix_memalign(void **memptr, size_t alignment, size_t size);  
void *aligned_alloc(size_t alignment, size_t size);  
void *valloc(size_t size);
```

```
#include <malloc.h>
```

```
void *memalign(size_t alignment, size_t size);  
void *pvalloc(size_t size);
```



Alignment check

```
template <typename T>
constexpr bool is_aligned (T* ptr, size_t align = alignof (T))
{
    return uintptr_t (ptr) % align == 0;
}
```


//...or, with checking whether align is power of 2:

```
template <typename T>
constexpr bool is_aligned (T* ptr, size_t align = alignof (T))
{
    if (! ((align != 0) && ! (align & (align - 1))))
        return false;

    return (uintptr_t (ptr) & (align - 1)) == 0;
}
```









**cache
line**

Align to cache lines



Align to cache lines



cacheline_size_t



SIMD


Single instruction, multiple data

SIMD Instructions

...

movups

mulps



CPU

SIMD Registers



xmm0



xmm1



xmm2

...



SIMD flavours

| SIMD flavour | Architecture | Registers | Register size |
|---------------------------|--------------|----------------------------|---------------|
| SSE (SSE2/3/4/4.1/4.2) | x86 | xmm0, ... xmm7 | 16 Bytes |
| AVX, AVX2 | x86 | ymm0, ... ymm15 | 32 Bytes |
| AVX512 | x86 | zmm0, ... zmm31 | 64 Bytes |
| NEON | ARM | q0, ... q15 d0, ... d31 | 16 Bytes |

SIMD and you

-03

```
for (size_t i = 0; i < size; ++i)
    a[i] *= f;
```

Auto-vectoriser

*JUCE, Boost.SIMD,
veclib, ...*

```
Vec4<float> a;
a *= f;
```

SIMD library

```
#include <emmintrin.h>

__m128 xmm0 = _mm_load_ps (a);
__m128 xmm1 = _mm_load1_ps (f);

xmm0 = _mm_mul_ps (xmm0, xmm1);
_mm_store_ps (xmm0, a);
```

SSE intrinsics

Inline assembly

```
__asm
{
    mov     eax, a
    mov     ebx, f
    movups xmm0, [eax]
    movups xmm1, [ebx]
    mulps  xmm1, xmm0
    ...
}
```

SIMD and you

Auto-vectoriser

SIMD library

SSE intrinsics

Inline assembly

easy

portable



hard

platform-specific

SIMD and you

Auto-vectoriser

SIMD library

SSE intrinsics

Inline assembly

easy

portable



hard

*platform-
specific*



```
void multiply (float* buffer, size_t size, float factor)
{
    for (size_t i = 0; i < size; ++i)
        buffer[i] *= factor;
}
```

```
void multiply (float* buffer, size_t size, float factor)
{
    for (size_t i = 0; i < size; ++i)
        buffer[i] *= factor;
}
```



clang++ -O3 -S

?

```

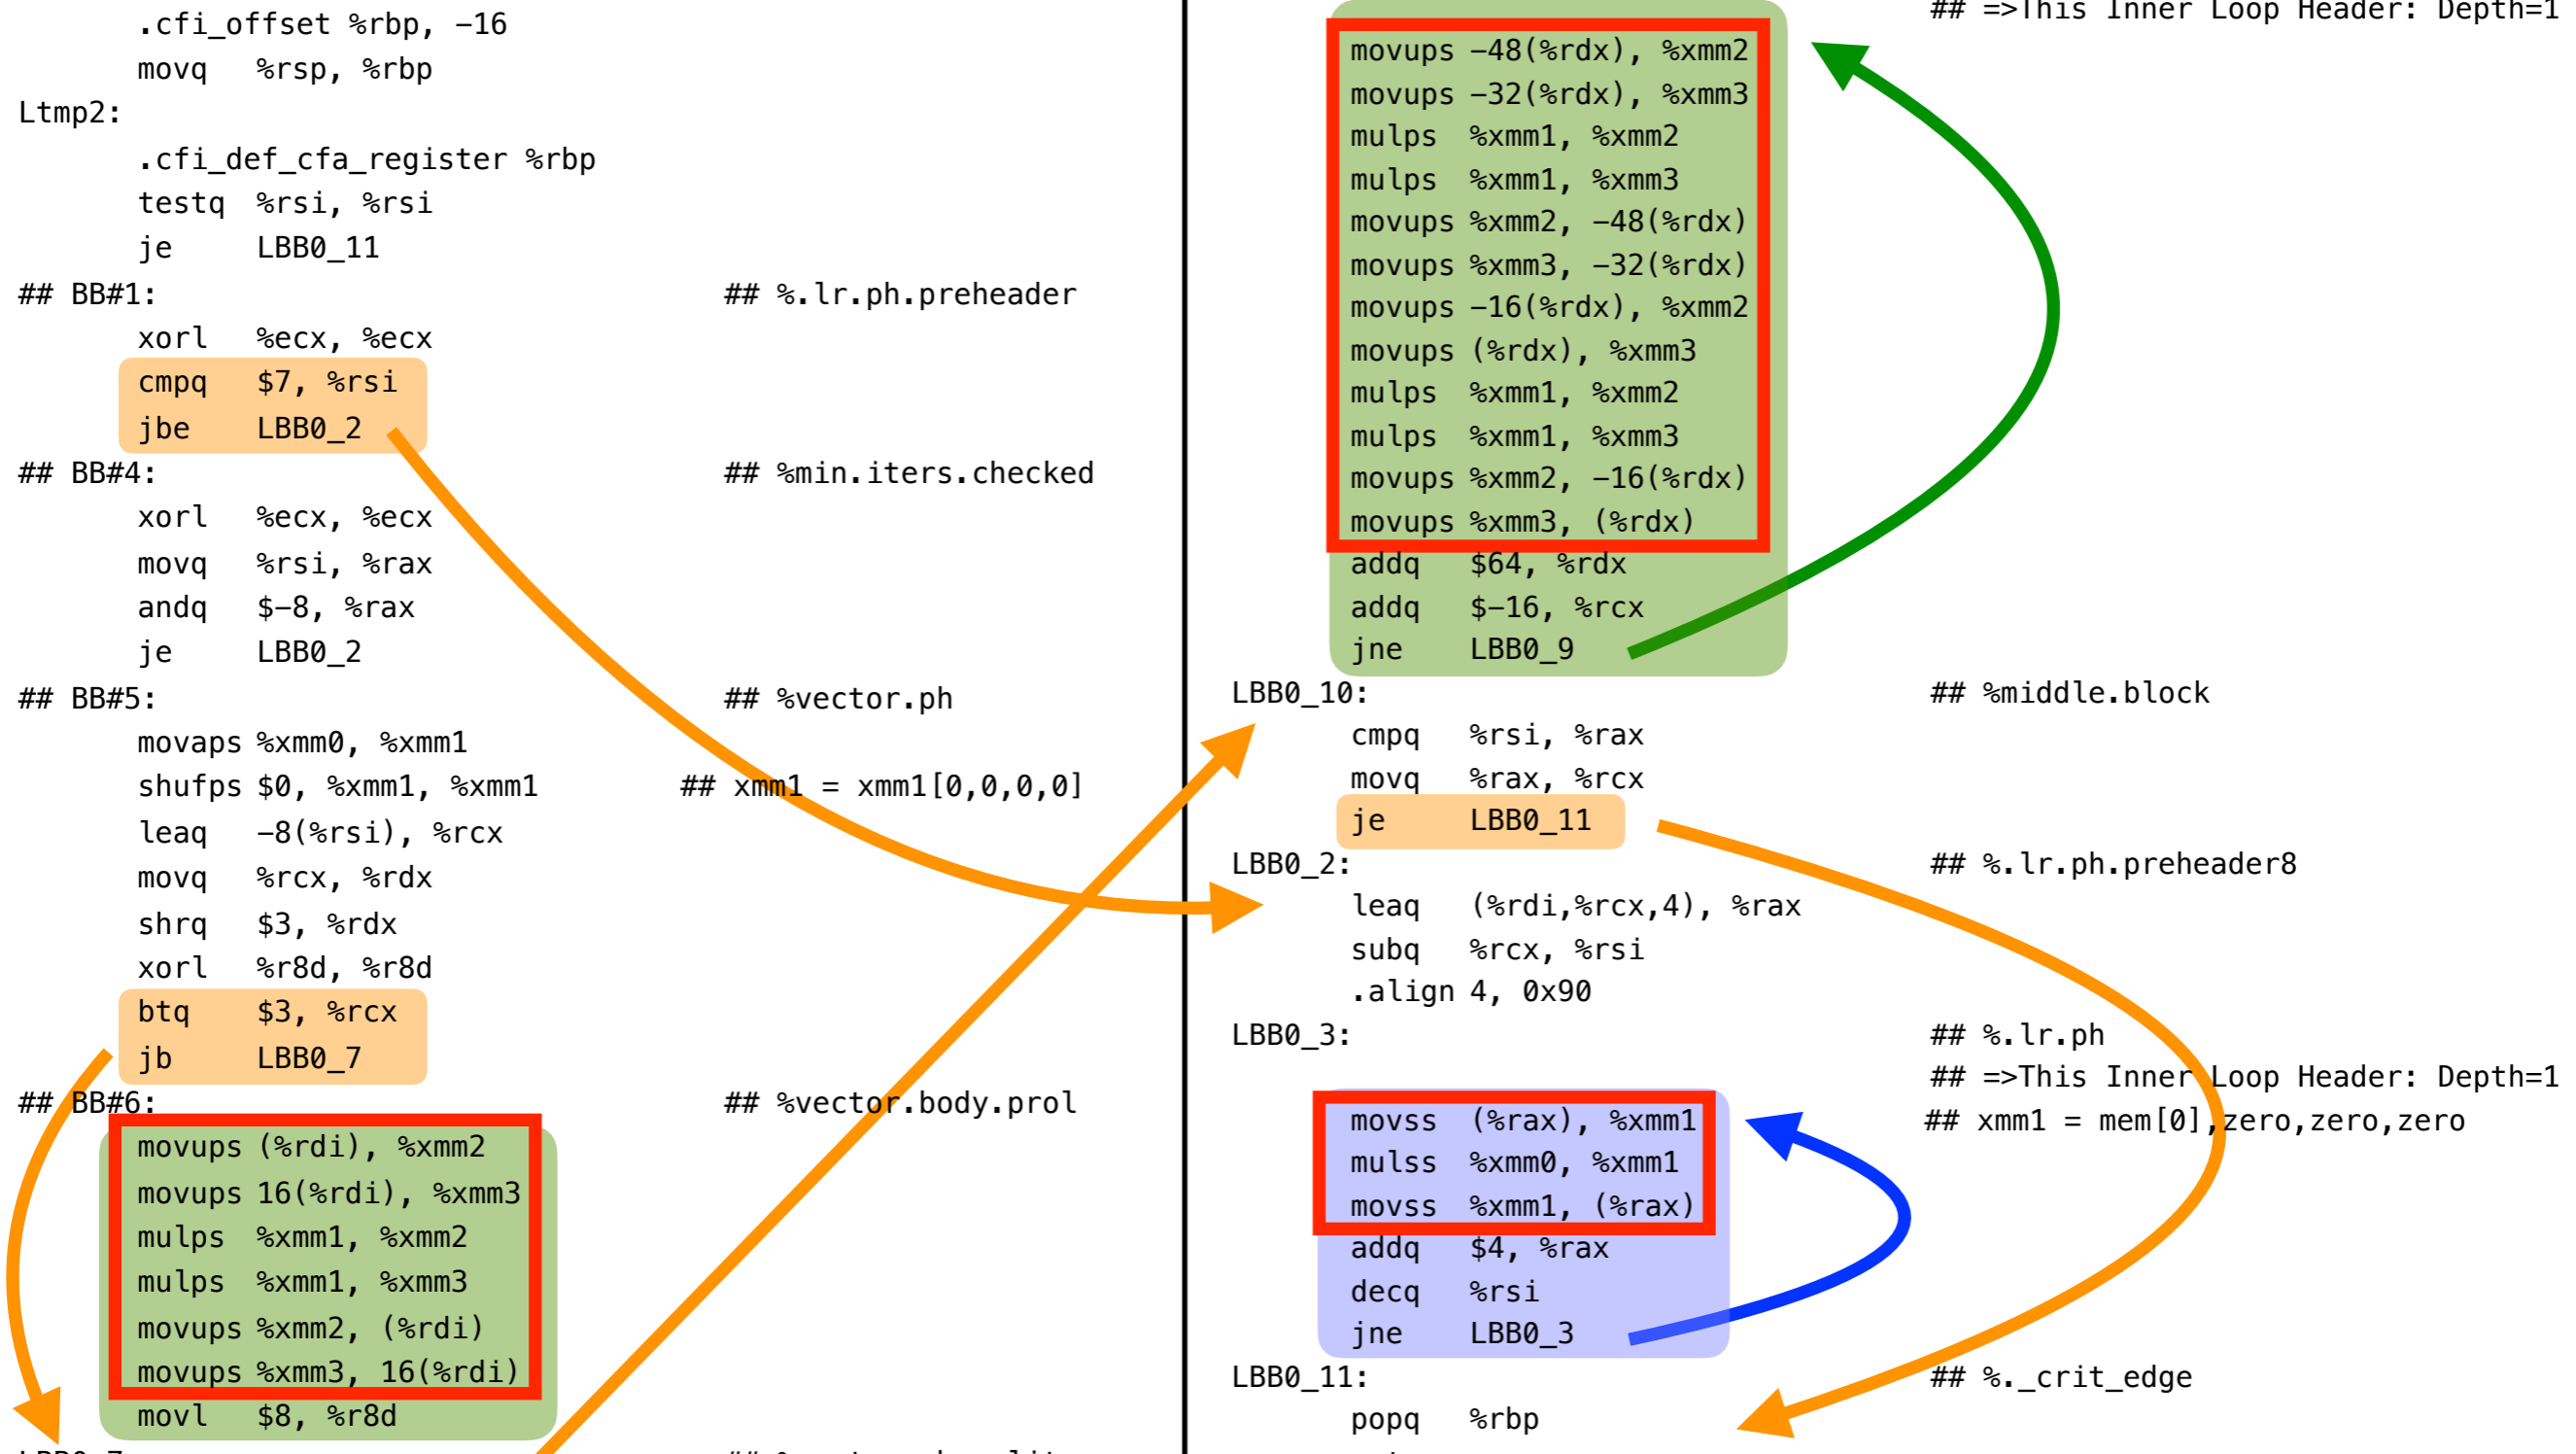
__Z8multiplyPfmf:
    .cfi_startproc
## BB#0:
    pushq %rbp
Ltmp0:
    .cfi_def_cfa_offset 16
Ltmp1:
    .cfi_offset %rbp, -16
    movq %rsp, %rbp
Ltmp2:
    .cfi_def_cfa_register %rbp
    testq %rsi, %rsi
    je    LBB0_11
## BB#1:
    xorl %ecx, %ecx
    cmpq $7, %rsi
    jbe  LBB0_2
## BB#4:
    xorl %ecx, %ecx
    movq %rsi, %rax
    andq $-8, %rax
    je    LBB0_2
## BB#5:
    movaps %xmm0, %xmm1
    shufps $0, %xmm1, %xmm1
    leaq -8(%rsi), %rcx
    movq %rcx, %rdx
    shrq $3, %rdx
    xorl %r8d, %r8d
    btq $3, %rcx
    jb   LBB0_7
## BB#6:
    movups (%rdi), %xmm2
    movups 16(%rdi), %xmm3
    mulps %xmm1, %xmm2
    mulps %xmm1, %xmm3
    movups %xmm2, (%rdi)
    movups %xmm3, 16(%rdi)
    movl $8, %r8d
LBB0_7:
    testq %rdx, %rdx
    je    LBB0_10
## @_Z8multiplyPfmf
## %.lr.ph.preheader
## %.min.iters.checked
## %vector.ph
## xmm1 = xmm1[0,0,0,0]
## %vector.body.prol
## %vector.ph.split

```

```

## BB#8:
    movq %rsi, %rcx
    andq $-8, %rcx
    subq %r8, %rcx
    leaq 48(%rdi,%r8,4), %rdx
    .align 4, 0x90
LBB0_9:
    movups -48(%rdx), %xmm2
    movups -32(%rdx), %xmm3
    mulps %xmm1, %xmm2
    mulps %xmm1, %xmm3
    movups %xmm2, -48(%rdx)
    movups %xmm3, -32(%rdx)
    movups -16(%rdx), %xmm2
    movups (%rdx), %xmm3
    mulps %xmm1, %xmm2
    mulps %xmm1, %xmm3
    movups %xmm2, -16(%rdx)
    movups %xmm3, (%rdx)
    addq $64, %rdx
    addq $-16, %rcx
    jne  LBB0_9
LBB0_10:
    cmpq %rsi, %rax
    movq %rax, %rcx
    je   LBB0_11
LBB0_2:
    leaq (%rdi,%rcx,4), %rax
    subq %rcx, %rsi
    .align 4, 0x90
LBB0_3:
    movss (%rax), %xmm1
    mulss %xmm0, %xmm1
    movss %xmm1, (%rax)
    addq $4, %rax
    decq %rsi
    jne  LBB0_3
LBB0_11:
    popq %rbp
    retq
.cfi_endproc
## %vector.ph.split.split
## %vector.body
## =>This Inner Loop Header: Depth=1
## %middle.block
## %.lr.ph.preheader8
## %.lr.ph
## =>This Inner Loop Header: Depth=1
## xmm1 = mem[0],zero,zero,zero
## %._crit_edge

```



```
void multiply (float* buffer, size_t size, float factor)
{
    for (size_t i = 0; i < size; ++i)
        buffer[i] *= factor;
}
```



```
template <size_t size>
void multiply (float* buffer, float factor)
{
    for (size_t i = 0; i < size; ++i)
        buffer[i] *= factor;
}
```

```
template void multiply<16384> (float* buffer, float factor);
```




clang++ -O3 -S



```

.cfi_startproc
## BB#0:                                     ##
%min.iters.checked
    pushq %rbp
Ltmp0:
    .cfi_def_cfa_offset 16
Ltmp1:
    .cfi_offset %rbp, -16
    movq %rsp, %rbp
Ltmp2:
    .cfi_def_cfa_register %rbp
    shufps    $0, %xmm0, %xmm0             ## xmm0 =
xmm0[0,0,0,0]
    xorl %eax, %eax
    .align    4, 0x90
LBB0_1:                                     ## %vector.body
                                     ## =>This Inner
Loop Header: Depth=1
movups    (%rdi,%rax,4), %xmm1
movups    16(%rdi,%rax,4), %xmm2
mulps %xmm0, %xmm1
mulps %xmm0, %xmm2
movups    %xmm1, (%rdi,%rax,4)
movups    %xmm2, 16(%rdi,%rax,4)
movups    32(%rdi,%rax,4), %xmm1
movups    48(%rdi,%rax,4), %xmm2
mulps %xmm0, %xmm1
mulps %xmm0, %xmm2
movups    %xmm1, 32(%rdi,%rax,4)
movups    %xmm2, 48(%rdi,%rax,4)
addq    $16, %rax
cmpq    $16384, %rax                       ## imm = 0x4000
jne    LBB0_1
## BB#2:                                     ## %middle.block
    popq %rbp
    retq
.cfi_endproc

```



```
void multiply (float* buffer, size_t size, float factor)
{
    for (size_t i = 0; i < size; ++i)
        buffer[i] *= factor;
}
```



```
template <size_t size>
void multiply (float* buffer, float factor)
{
    for (size_t i = 0; i < size; ++i)
        buffer[i] *= factor;
}
```

```
template void multiply<128> (float* buffer, float factor);
```



clang++ -03 -S



```

__Z8multiplyILm128EEvPff:
@_Z8multiplyILm128EEvPff
    .cfi_startproc
## BB#0:
    pushq %rbp
Ltmp0:
    .cfi_def_cfa_offset 16
Ltmp1:
    .cfi_offset %rbp, -16
    movq %rsp, %rbp
Ltmp2:
    .cfi_def_cfa_register %rbp
    shufps $0, %xmm0, %xmm0

```

```

    movups (%rdi), %xmm1
    movups 16(%rdi), %xmm2
    mulps %xmm0, %xmm1
    mulps %xmm0, %xmm2
    movups %xmm1, (%rdi)
    movups %xmm2, 16(%rdi)
    movups 32(%rdi), %xmm1
    movups 48(%rdi), %xmm2
    mulps %xmm0, %xmm1
    mulps %xmm0, %xmm2
    movups %xmm1, 32(%rdi)
    movups %xmm2, 48(%rdi)
    movups 64(%rdi), %xmm1
    movups 80(%rdi), %xmm2
    mulps %xmm0, %xmm1
    mulps %xmm0, %xmm2
    movups %xmm1, 64(%rdi)
    movups %xmm2, 80(%rdi)
    movups 96(%rdi), %xmm1
    movups 112(%rdi), %xmm2
    mulps %xmm0, %xmm1
    mulps %xmm0, %xmm2
    movups %xmm1, 96(%rdi)
    movups %xmm2, 112(%rdi)
    movups 128(%rdi), %xmm1

```

```

##
## %min.iters.checked

```

```

## xmm0 = xmm0[0,0,0,0]

```

```

    movups 144(%rdi), %xmm2
    mulps %xmm0, %xmm1
    mulps %xmm0, %xmm2
    movups %xmm1, 128(%rdi)
    movups %xmm2, 144(%rdi)
    movups 160(%rdi), %xmm1
    movups 176(%rdi), %xmm2
    mulps %xmm0, %xmm1
    mulps %xmm0, %xmm2
    movups %xmm1, 160(%rdi)
    movups %xmm2, 176(%rdi)
    movups 192(%rdi), %xmm1
    movups 208(%rdi), %xmm2
    mulps %xmm0, %xmm1
    mulps %xmm0, %xmm2
    movups %xmm1, 192(%rdi)
    movups %xmm2, 208(%rdi)
    movups 224(%rdi), %xmm1
    movups 240(%rdi), %xmm2
    mulps %xmm0, %xmm1
    mulps %xmm0, %xmm2
    movups %xmm1, 224(%rdi)
    movups %xmm2, 240(%rdi)
    movups 256(%rdi), %xmm1
    movups 272(%rdi), %xmm2
    mulps %xmm0, %xmm1
    mulps %xmm0, %xmm2
    movups %xmm1, 256(%rdi)
    movups %xmm2, 272(%rdi)
    movups 288(%rdi), %xmm1
    movups 304(%rdi), %xmm2
    mulps %xmm0, %xmm1
    mulps %xmm0, %xmm2
    movups %xmm1, 288(%rdi)
    movups %xmm2, 304(%rdi)
    movups 320(%rdi), %xmm1
    movups 336(%rdi), %xmm2

```


```

    mulps %xmm0, %xmm1
    mulps %xmm0, %xmm2
    movups %xmm1, 320(%rdi)
    movups %xmm2, 336(%rdi)
    movups 352(%rdi), %xmm1
    movups 368(%rdi), %xmm2
    mulps %xmm0, %xmm1
    mulps %xmm0, %xmm2
    movups %xmm1, 352(%rdi)
    movups %xmm2, 368(%rdi)
    movups 384(%rdi), %xmm1
    movups 400(%rdi), %xmm2
    mulps %xmm0, %xmm1
    mulps %xmm0, %xmm2
    movups %xmm1, 384(%rdi)
    movups %xmm2, 400(%rdi)
    movups 416(%rdi), %xmm1
    movups 432(%rdi), %xmm2
    mulps %xmm0, %xmm1
    mulps %xmm0, %xmm2
    movups %xmm1, 416(%rdi)
    movups %xmm2, 432(%rdi)
    movups 448(%rdi), %xmm1
    movups 464(%rdi), %xmm2
    mulps %xmm0, %xmm1
    mulps %xmm0, %xmm2
    movups %xmm1, 448(%rdi)
    movups %xmm2, 464(%rdi)
    movups 480(%rdi), %xmm1
    movups 496(%rdi), %xmm2
    mulps %xmm0, %xmm1
    mulps %xmm0, %xmm2
    movups %xmm1, 480(%rdi)
    movups %xmm2, 496(%rdi)
    popq %rbp
    retq
    .cfi_endproc

```



```
void multiply (float* buffer, size_t size, float factor)
{
    for (size_t i = 0; i < size; ++i)
        buffer[i] *= factor;
}
```



?

__Z8multiplyILm16384EEvPff:

LFB1:

```
movq %rdi, %rax
andl $15, %eax
shrq $2, %rax
negq %rax
andl $3, %eax
je L8
```

```
movss (%rdi), %xmm1
cmpq $1, %rax
mulss %xmm0, %xmm1
movss %xmm1, (%rdi)
je L9
```

```
movss 4(%rdi), %xmm1
cmpq $2, %rax
mulss %xmm0, %xmm1
movss %xmm1, 4(%rdi)
je L10
```

```
movss 8(%rdi), %xmm1
movl $16381, %r9d
movl $3, %r8d
mulss %xmm0, %xmm1
movss %xmm1, 8(%rdi)
```

L3:

```
movl $16384, %r10d
movl $16380, %r11d
movl $4095, %esi
subq %rax, %r10
```

L2:

```
leaq (%rdi,%rax,4), %rcx
movaps %xmm0, %xmm2
xorl %eax, %eax
xorl %edx, %edx
shufps $0, %xmm2, %xmm2
.align 4,0x90
```

L5:

```
movaps (%rcx,%rax), %xmm1
addq $1, %rdx
mulps %xmm2, %xmm1
movaps %xmm1, (%rcx,%rax)
addq $16, %rax
cmpq %rsi, %rdx
jb L5
```

```
leaq (%r8,%r11), %rax
movq %r9, %rdx
subq %r11, %rdx
cmpq %r11, %r10
je L1
leaq (%rdi,%rax,4), %rcx
cmpq $1, %rdx
```

```
movss (%rcx), %xmm1
mulss %xmm0, %xmm1
movss %xmm1, (%rcx)
```

```
leaq 1(%rax), %rcx
je L1
leaq (%rdi,%rcx,4), %rcx
addq $2, %rax
cmpq $2, %rdx
```

```
movss (%rcx), %xmm1
mulss %xmm0, %xmm1
movss %xmm1, (%rcx)
```

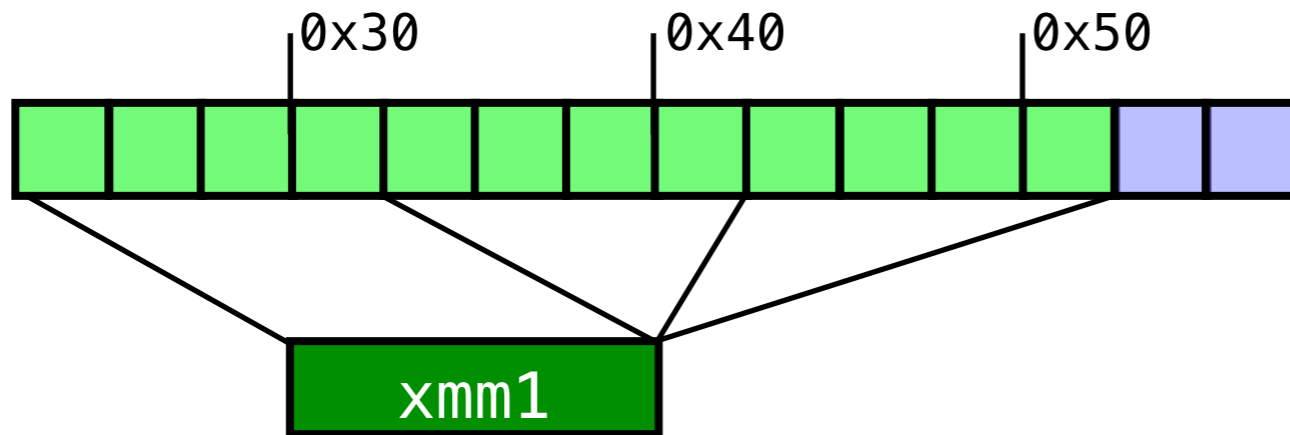
```
je L1
leaq (%rdi,%rax,4), %rax
mulss (%rax), %xmm0
movss %xmm0, (%rax)
```

```
ret
.align 4,0x90
```

L1:


```
ret
```





movups

(unaligned load/store from/to SSE registers)



movaps

(aligned load/store from/to SSE registers)

multiplyAdd benchmark (SSE)

```
void multiplyAdd (float* buffer1, float* buffer2,  
                 float factor, size_t size)  
{  
    for (size_t i = 0; i < size; ++i)  
        buffer1[i] += buffer2[i] * factor;  
}
```

multiplyAdd benchmark (SSE)

```
void multiplyAdd (float* buffer1, float* buffer2, float factor, std::size_t size)
{
    __m128 f = _mm_load1_ps (&factor);

    for (std::size_t i = 0; i < size; i+=4)
    {
        __m128 b1 = _mm_loadu_ps (buffer1 + i);           // or _mm_load_ps for aligned load
        __m128 b2 = _mm_loadu_ps (buffer2 + i);

        b2 = _mm_mul_ps (b2, f);
        b1 = _mm_add_ps (b1, b2);

        _mm_storeu_ps (buffer1 + i, b1);                // or _mm_store_ps for aligned load
    }
}

int main()
{
    const int size = 16384;

    auto* b1 = new float[size];
    auto* b2 = new float[size];

    // Warm the cache.
    std::fill (b2, b2 + size, 0.2f);
    std::fill (b1, b1 + size, 1.0f);

    BENCHMARK_START(loop, 100000)

    multiplyAdd (b1, b2, 0.00001f, size);                // or multiplyAdd (b1 + 1, b2 + 2, 0.00001f, size - 4);
                                                         // for non-SSE-aligned data

    BENCHMARK_STOP(loop)

    std::cout << b1[0] << std::endl;
}
```

multiplyAdd benchmark (NEON)

```
void multiplyAdd (float* buffer1, float* buffer2, float factor, std::size_t size)
{
    float32x4_t f = vld1q_f32 (&factor);

    for (std::size_t i = 0; i < size; i+=4)
    {
        float32x4_t b1 = vld1q_f32 (buffer1 + i);
        float32x4_t b2 = vld1q_f32 (buffer2 + i);

        b2 = vmulq_f32 (b2, f);
        b1 = vaddq_f32 (b1, b2);

        vst1q_f32 (buffer1 + i, b1);
    }
}

int main()
{
    const int size = 16384;

    auto* b1 = new float[size];
    auto* b2 = new float[size];

    // Warm the cache.
    std::fill (b2, b2 + size, 0.2f);
    std::fill (b1, b1 + size, 1.0f);

    BENCHMARK_START(loop, 100000)

    multiplyAdd (b1, b2, 0.00001f, size);           // or multiplyAdd (b1 + 1, b2 + 2, 0.00001f, size - 4);
                                                    // for non-SSE-aligned data

    BENCHMARK_STOP(loop)

    std::cout << b1[0] << std::endl;
}
```

Cost of unaligned move for SSE/NEON registers

(using the simple multiplyAdd benchmark)



src



* * * * *

kernel



+
||

dest



Convolution

src



* * * * *

kernel



+
||

dest




Convolution

src




* * * * *

kernel



+
||

dest

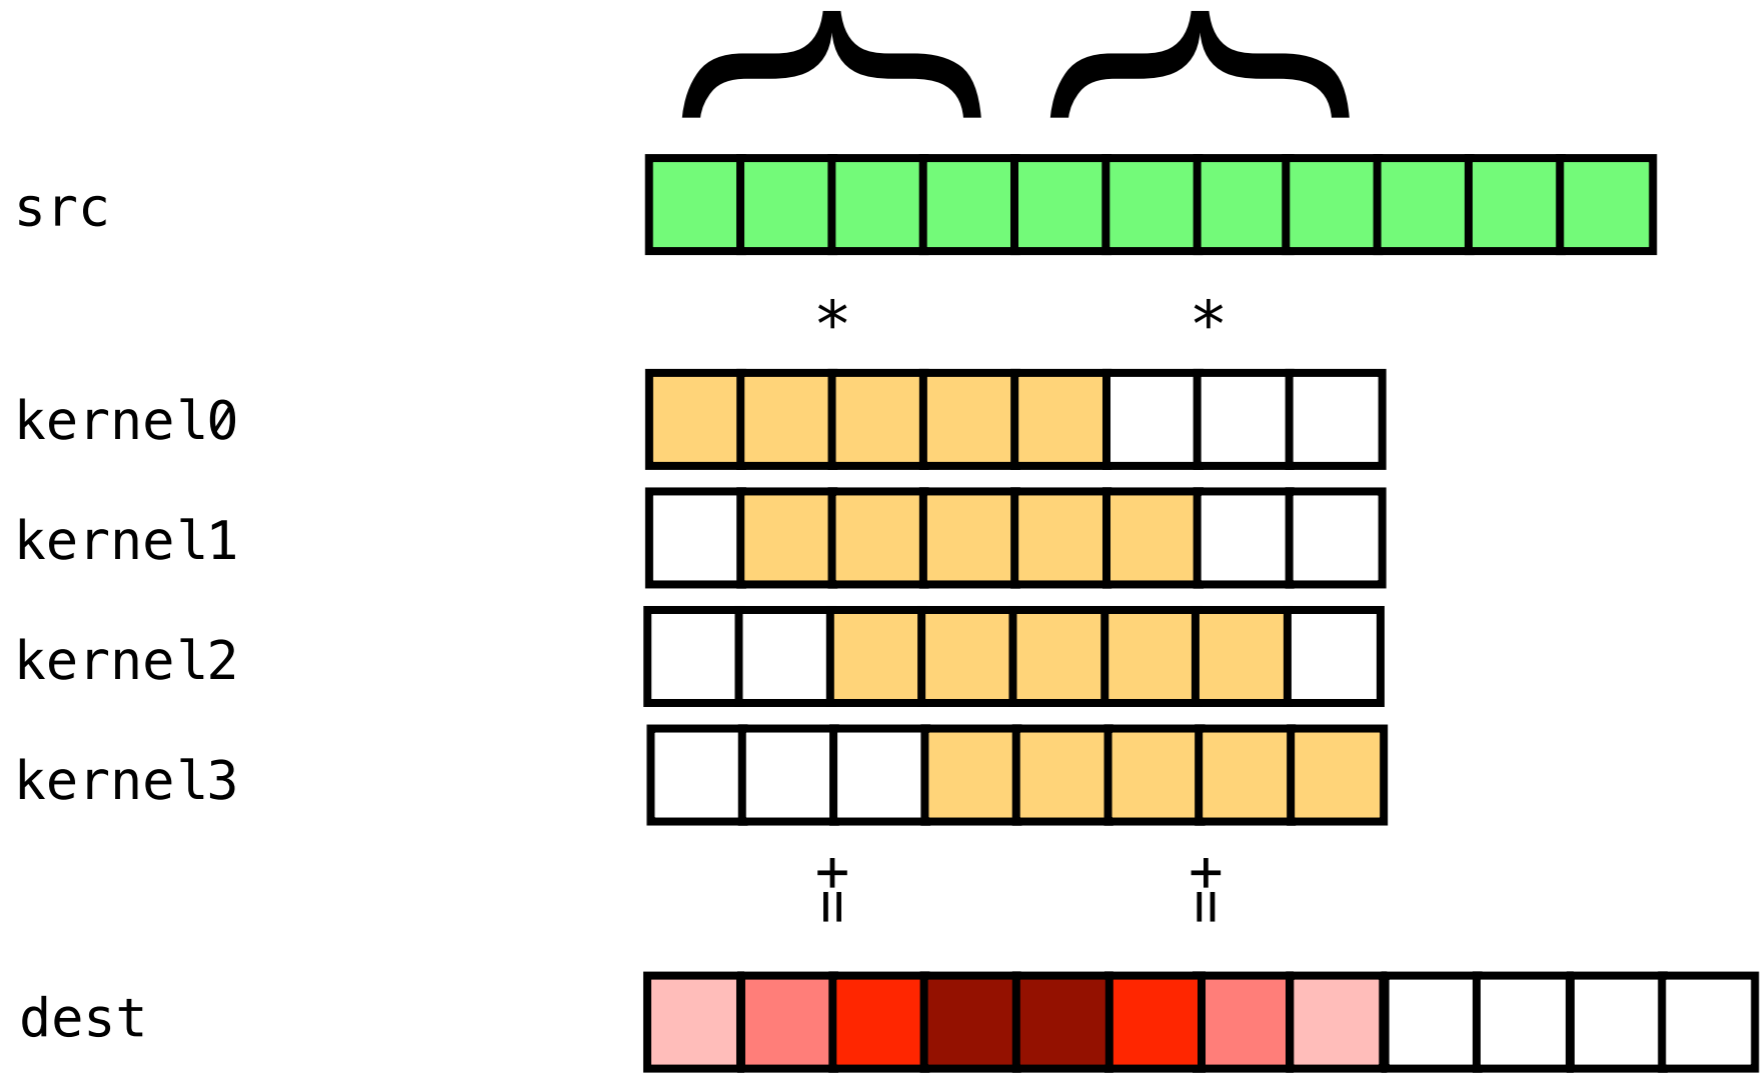


```
const float kernel[128];
float* src = new float[65536];
float* dest = new float[65536];

for (int i = 0; i < 65536 - 128; ++i)
{
    float sum = 0.0f;

    for (int j = 0; j < 128; ++j)
        sum += src [i + j] * kernel[j];

    dest[i] = sum;
}
```



```
void multiply (float* buffer, size_t size, float factor)
{
    // how to tell the compiler/optimiser at this point
    // that float* buffer is aligned to 16 bytes??

    for (size_t i = 0; i < size; ++i)
        buffer[i] *= factor;
}
```


// ...would be nice to have something like this:

```
template <typename T, size_t Align>
class aligned_ptr<T>;
```

// ...and then write something like this:

```
void multiply (aligned_ptr<float, 16> buffer, ...);
```





denormals

Crunch some floating point numbers!

```
float crunchSomeNumbers (float x)
{
    float sum = 0.0f;

    BENCHMARK_START(crunch, 100000)
    {
        for (int i = 0; i < 10000; ++i)
            x *= 0.999f;

        sum += x;
    }
    BENCHMARK_STOP(crunch)

    return sum;
}
```

x = 1.0;


x = 1.0 / 0.0;

x = 0.0 / 0.0;


x = -0.0;

x = 1e-39;

Number crunching benchmark



Denormals



- Exponent = 1111111:
 - fraction is zero \rightarrow +inf, -inf
 - fraction is non-zero \rightarrow NaNs
- Exponent between 1111110 and 00000001:
 - “normal” floats with implicit leading 1.
- Exponent = 0000000:
 - fraction is zero \rightarrow +0, -0
 - fraction is non-zero \rightarrow denormals

Check denormals support

```
std::numeric_limits<float>::min();           // 1.17549e-38
std::numeric_limits<float>::denorm_min();    // 1.4013e-45
std::numeric_limits<float>::has_denorm;     // std::denorm_present
```


Flush denormals to zero

```
#include <float.h>
```

```
// Flush denormals to zero, both operands and results  
_controlfp_s (nullptr, _DN_FLUSH, _MCW_DN);
```

```
// Put denormal handling back to normal.  
_controlfp_s (nullptr, _DN_SAVE, _MCW_DN);
```



The image features a dark, star-filled night sky as a background. The stars vary in size and color, including white, yellow, orange, and blue. The text "Thank you!" is centered in a bold, yellow, sans-serif font.

Thank you!