

Project Blizzard

Safe manual memory management

Alexandru Jercaianu
University POLITEHNICA of Bucharest
alex.jercaianu@gmail.com

DConf 2018
Munich, May 2-5, 2018

Manual memory management

High performance

Deterministic lifetimes

Garbage collector

Safety

Increase in productivity

Vulnerabilities

Use after free


Buffer overflows

Double free

Undefined behavior

```
1 struct Point
2 {
3     int x;
4     int y;
5 }
6
7 struct User
8 {
9     char[] name;
10    int age;
11 }
12
13 Point* p = Mallocator.instance.make!Point();
14 // do stuff with point
15 Mallocator.instance.dispose(p);
16
17 User* u = Mallocator.instance.make!User();
18 // ...
19
20 p.x = 100;
21 u.name[0] = 'a';
22
```





Blizzard allocator

Usable in safe code

Mitigate dangling pointers

High performance

Coexist with other allocators

Related work

- Project Snowflake: Non-blocking safe manual memory management in .NET

(<https://www.microsoft.com/en-us/research/publication/project-snowflake-non-blocking-safe-manual-memory-management-net/>)

- Simple, Fast and Safe Manual Memory Management

(<https://www.microsoft.com/en-us/research/publication/simple-fast-safe-manual-memory-management/>)

Allocators

```
1 struct MyAllocator
2 {
3     void[] allocate(size_t s);
4     bool deallocate(void[] b);
5
6     // Optional
7     void[] alignedAllocate(size_t s, uint a);
8     bool reallocate(ref void[] b, size_t newSize);
9     bool expand(ref void[] b, size_t delta);
10
11     // Many other primitives ...
12 }|
13
```

Allocator examples

Region

BitmappedBlock

GCAllocator

Mallocator

Building blocks

```
1 alias MyAllocator = Segregator!(  
2     64, BitmappedBlock!(64, 8),  
3     128, BitmappedBlock!(128, 16),  
4     Mallocator  
5 );
```

```
1 struct Point
2 {
3     int x;
4     int y;
5 }
6
7 struct User
8 {
9     char[] name;
10    int age;
11 }
12
13 Point* p = Mallocator.instance.make!Point();
14 // do stuff with point
15 Mallocator.instance.dispose(p);
16
17 User* u = Mallocator.instance.make!User();
18 // ...
19
20 p.x = 100;
21 u.name[0] = 'a';
22
```

Problem

Reusing memory is not safe

Problem

Reusing memory is not safe

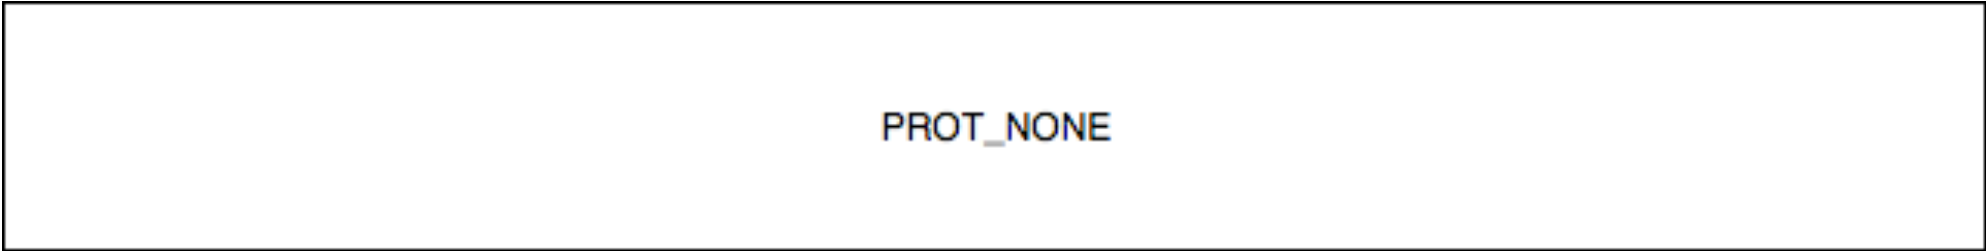
Solution

Always allocate at increasing addresses

Great for 64bit

Ascending Page Allocator


Map a large chunk of virtual memory with no permissions



PROT_NONE

Ascending Page Allocator

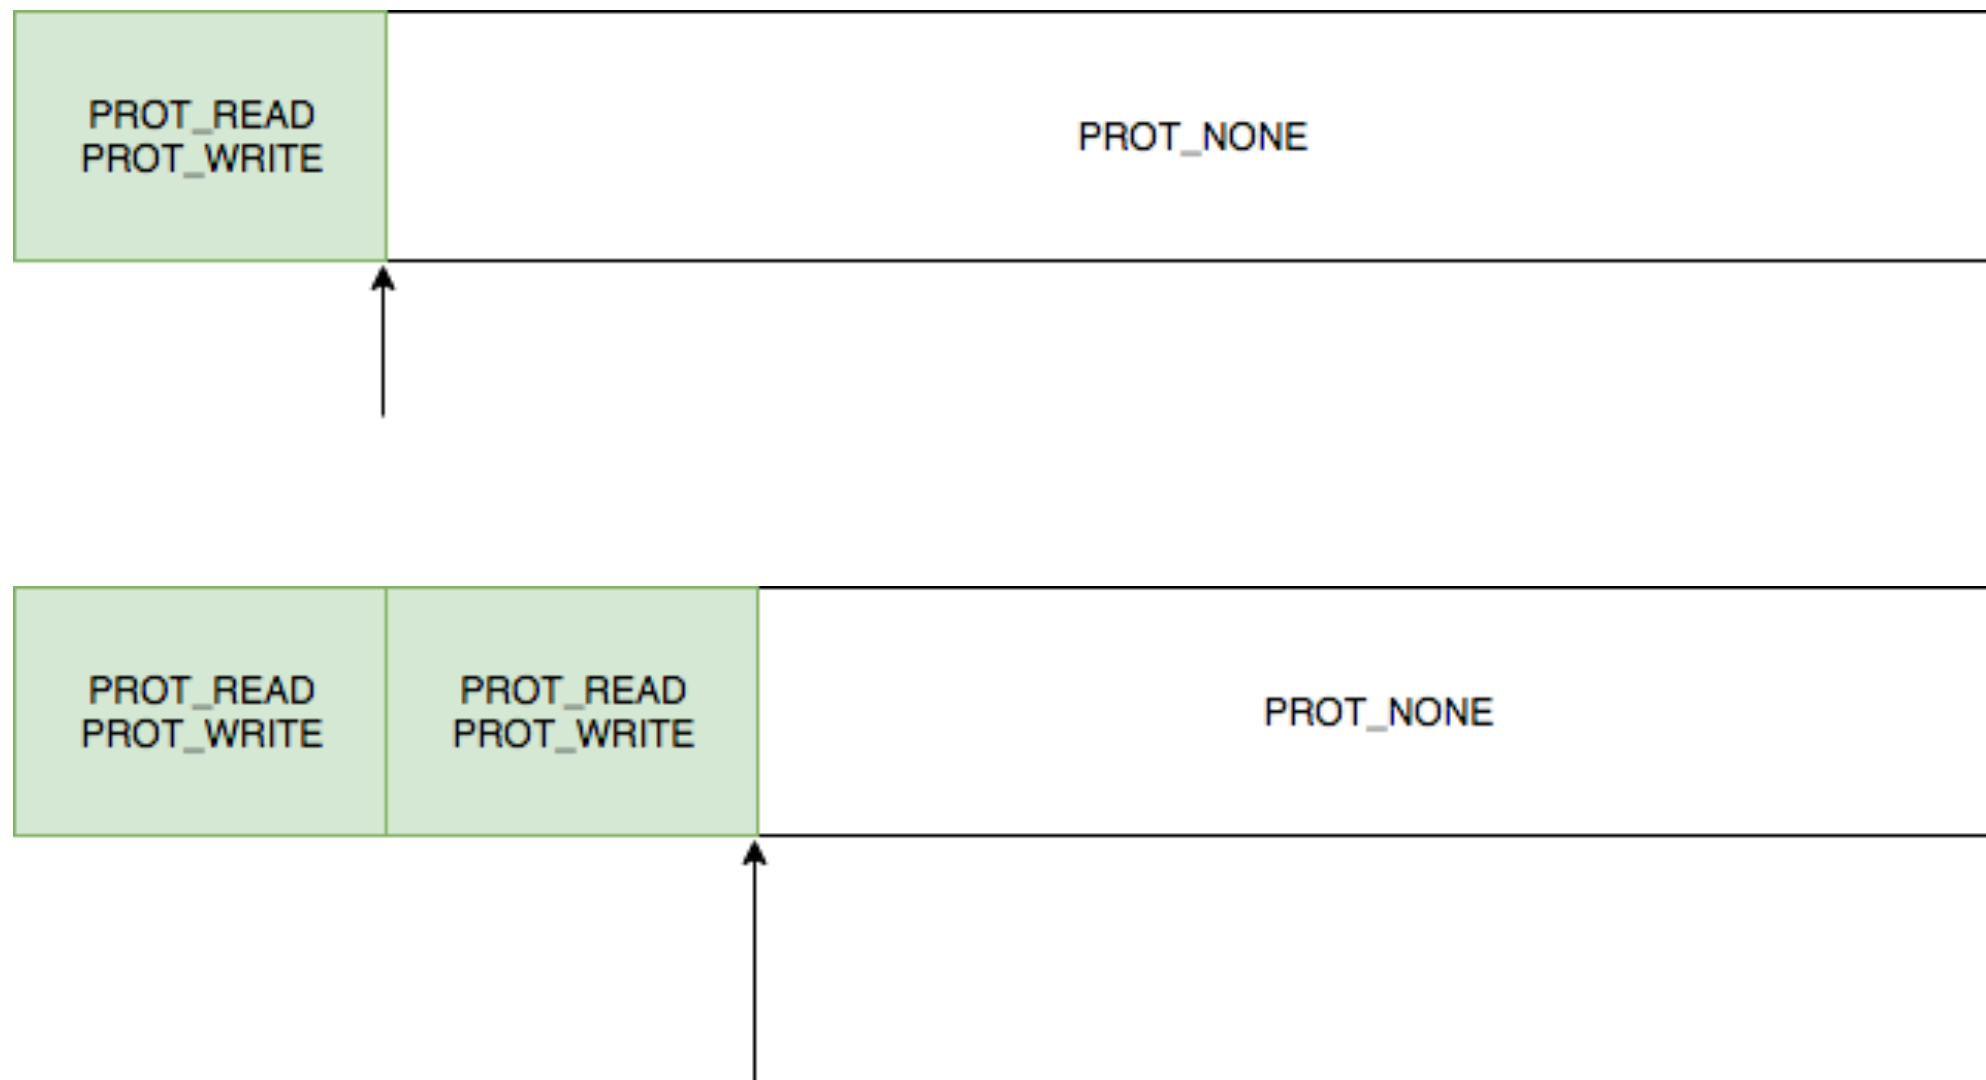
Map a large chunk of virtual memory with no permissions



```
1 // Posix
2 mmap(... PROT_NONE, MAP_PRIVATE | MAP_ANONYMOUS ...).
3
4 // Windows
5 VirtualAlloc(... MEM_RESERVE, PAGE_NOACCESS ...)
```

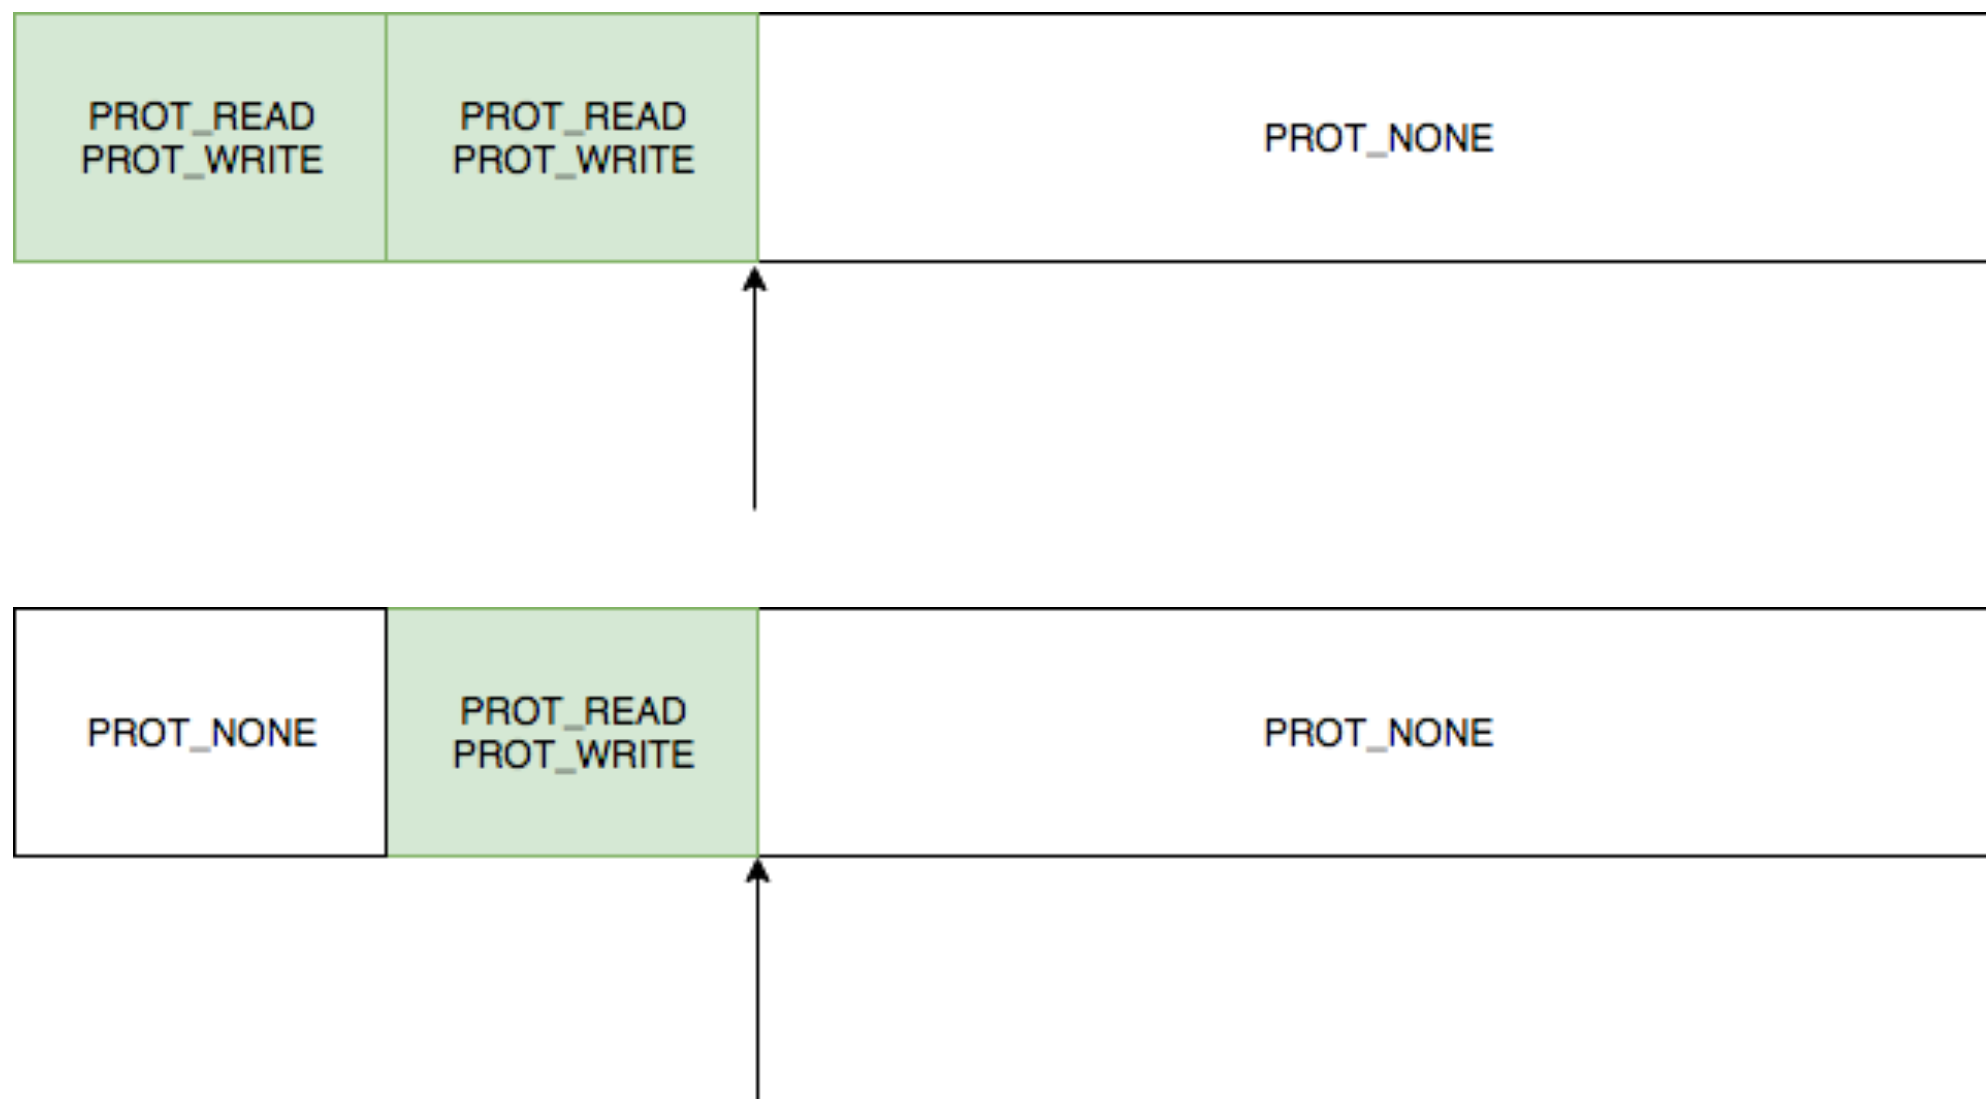

Ascending Page Allocator

Each allocation advances a pointer and sets read/write permissions



Ascending Page Allocator

Reclaim physical pages
Remove read/write permissions



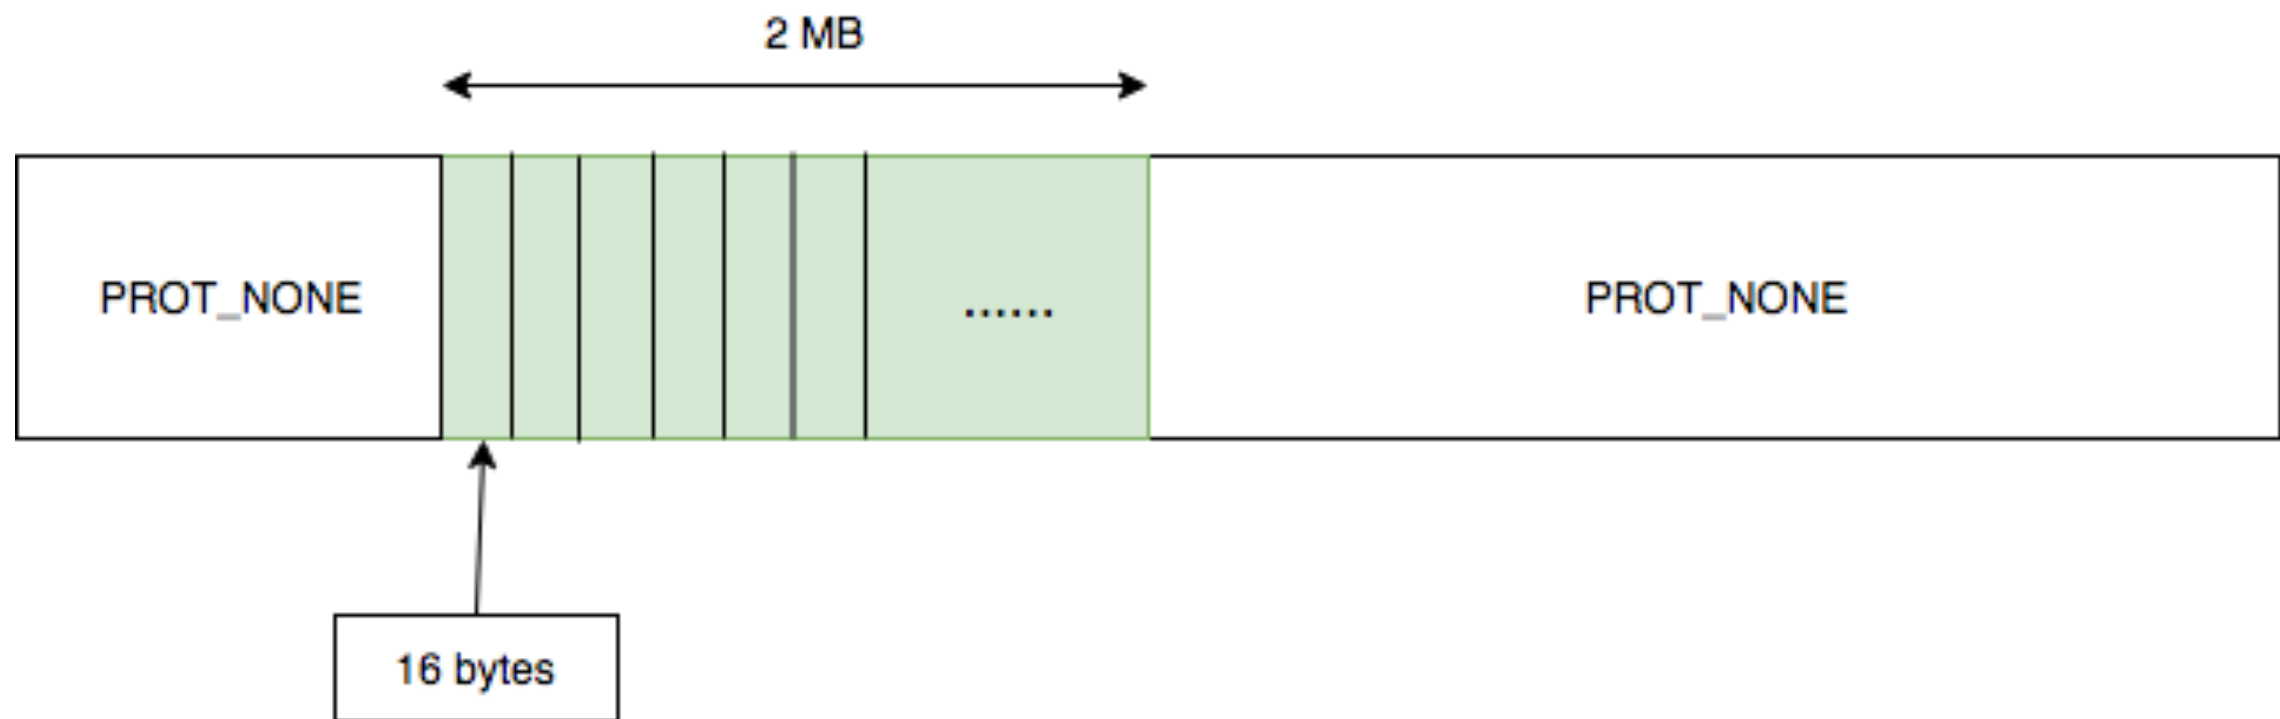
Ascending Page Allocator

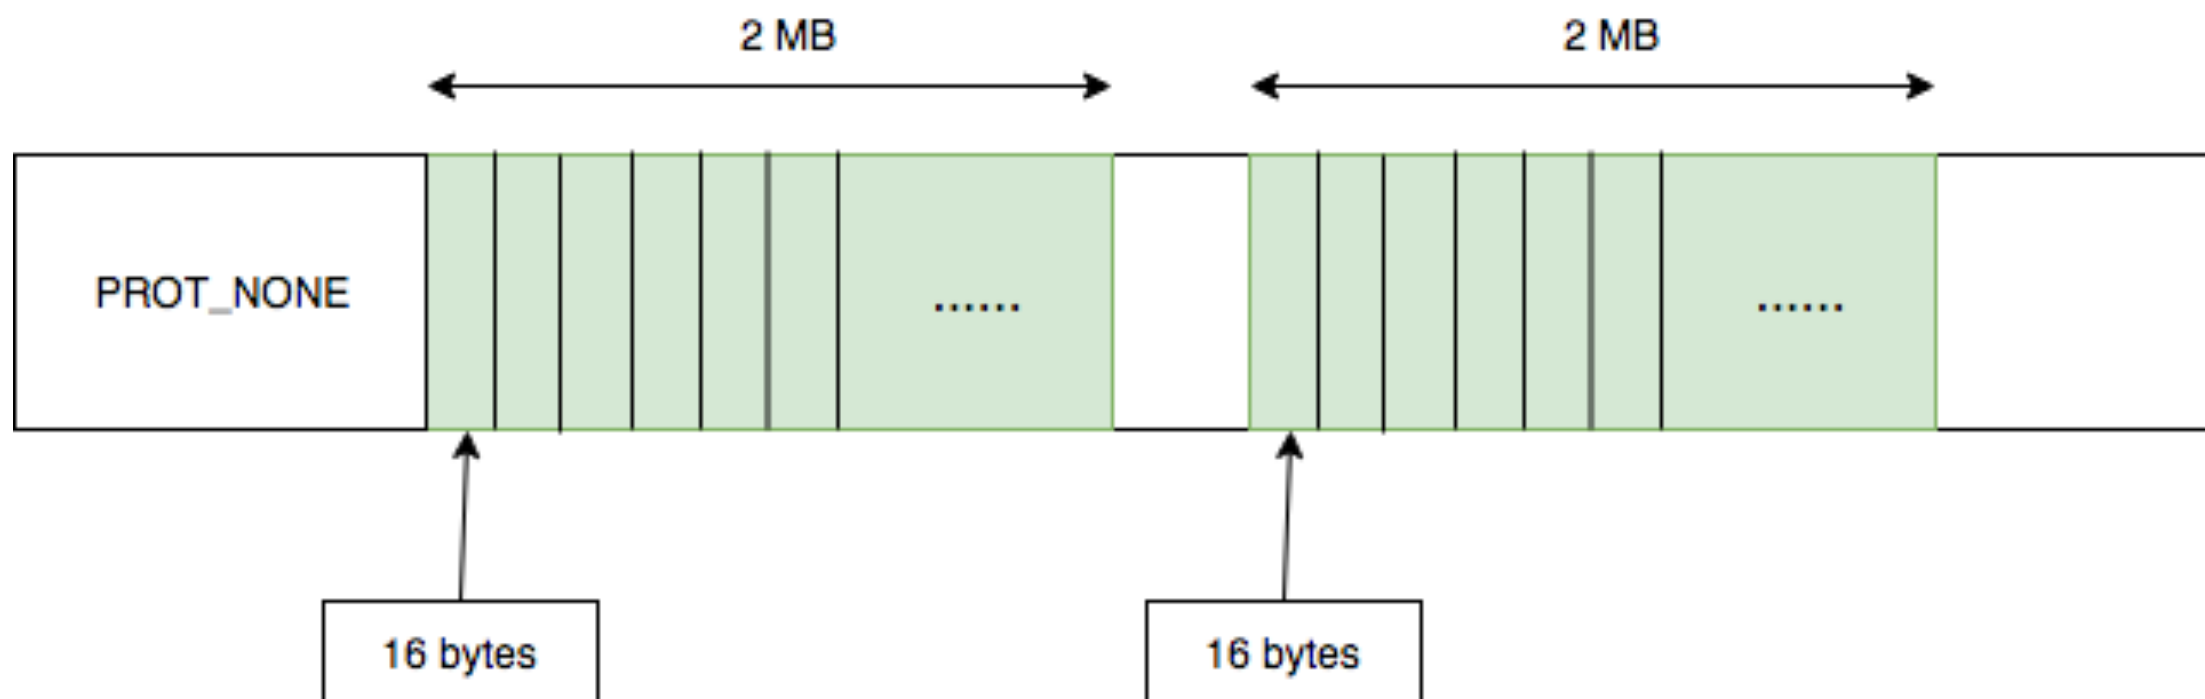
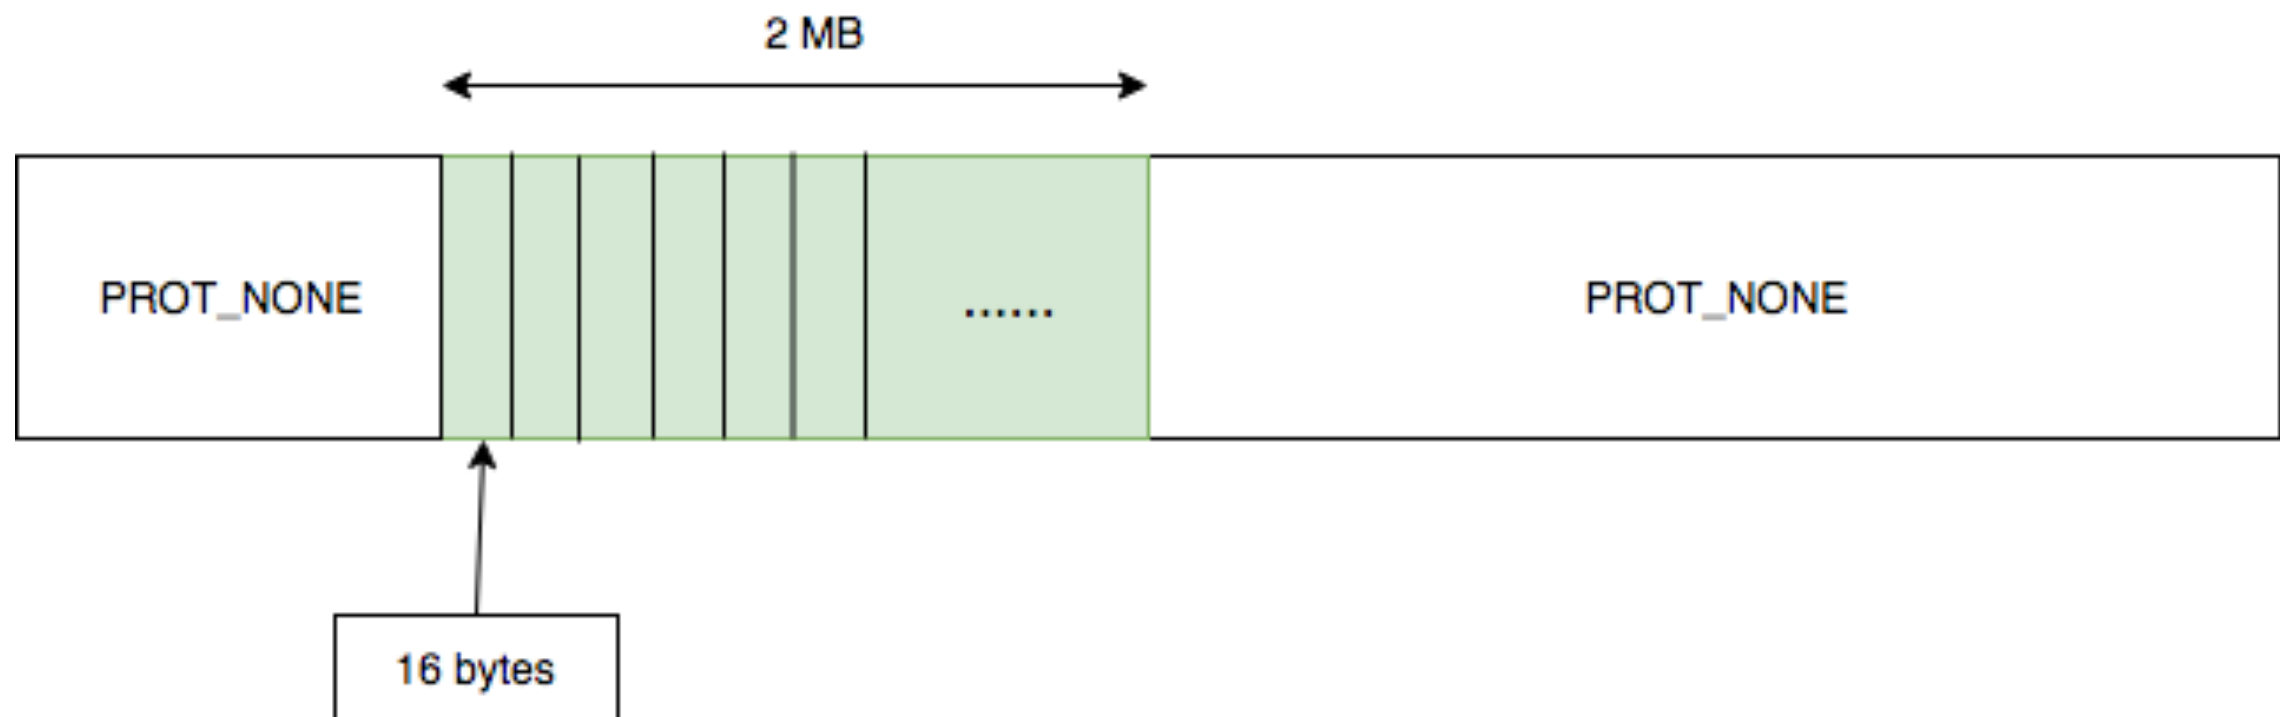
Safe

Page granularity

Great for huge allocations

High fragmentation



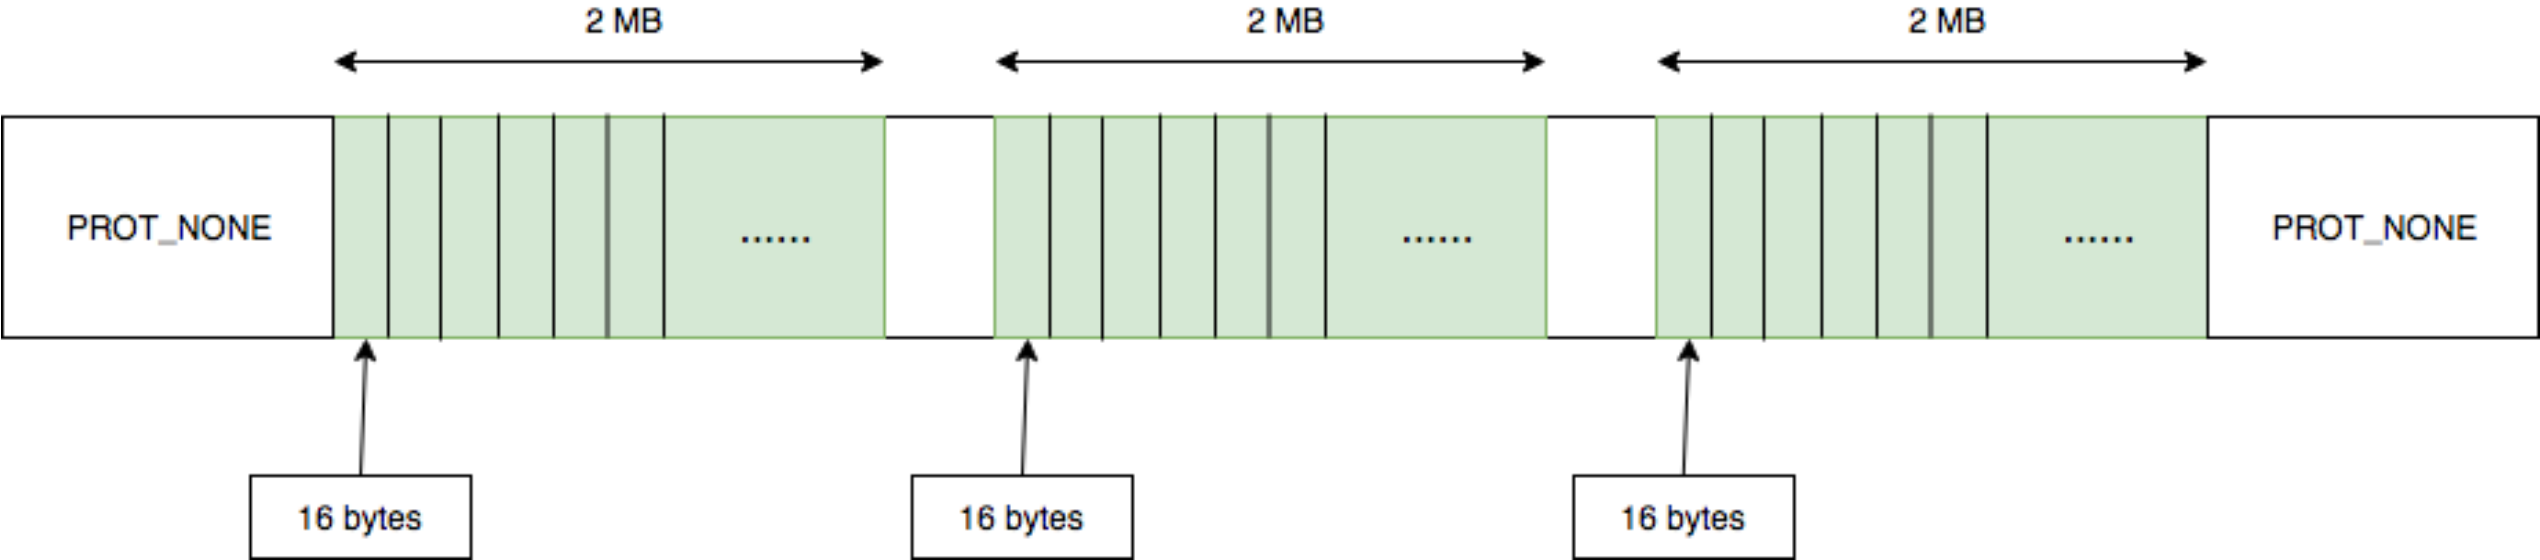


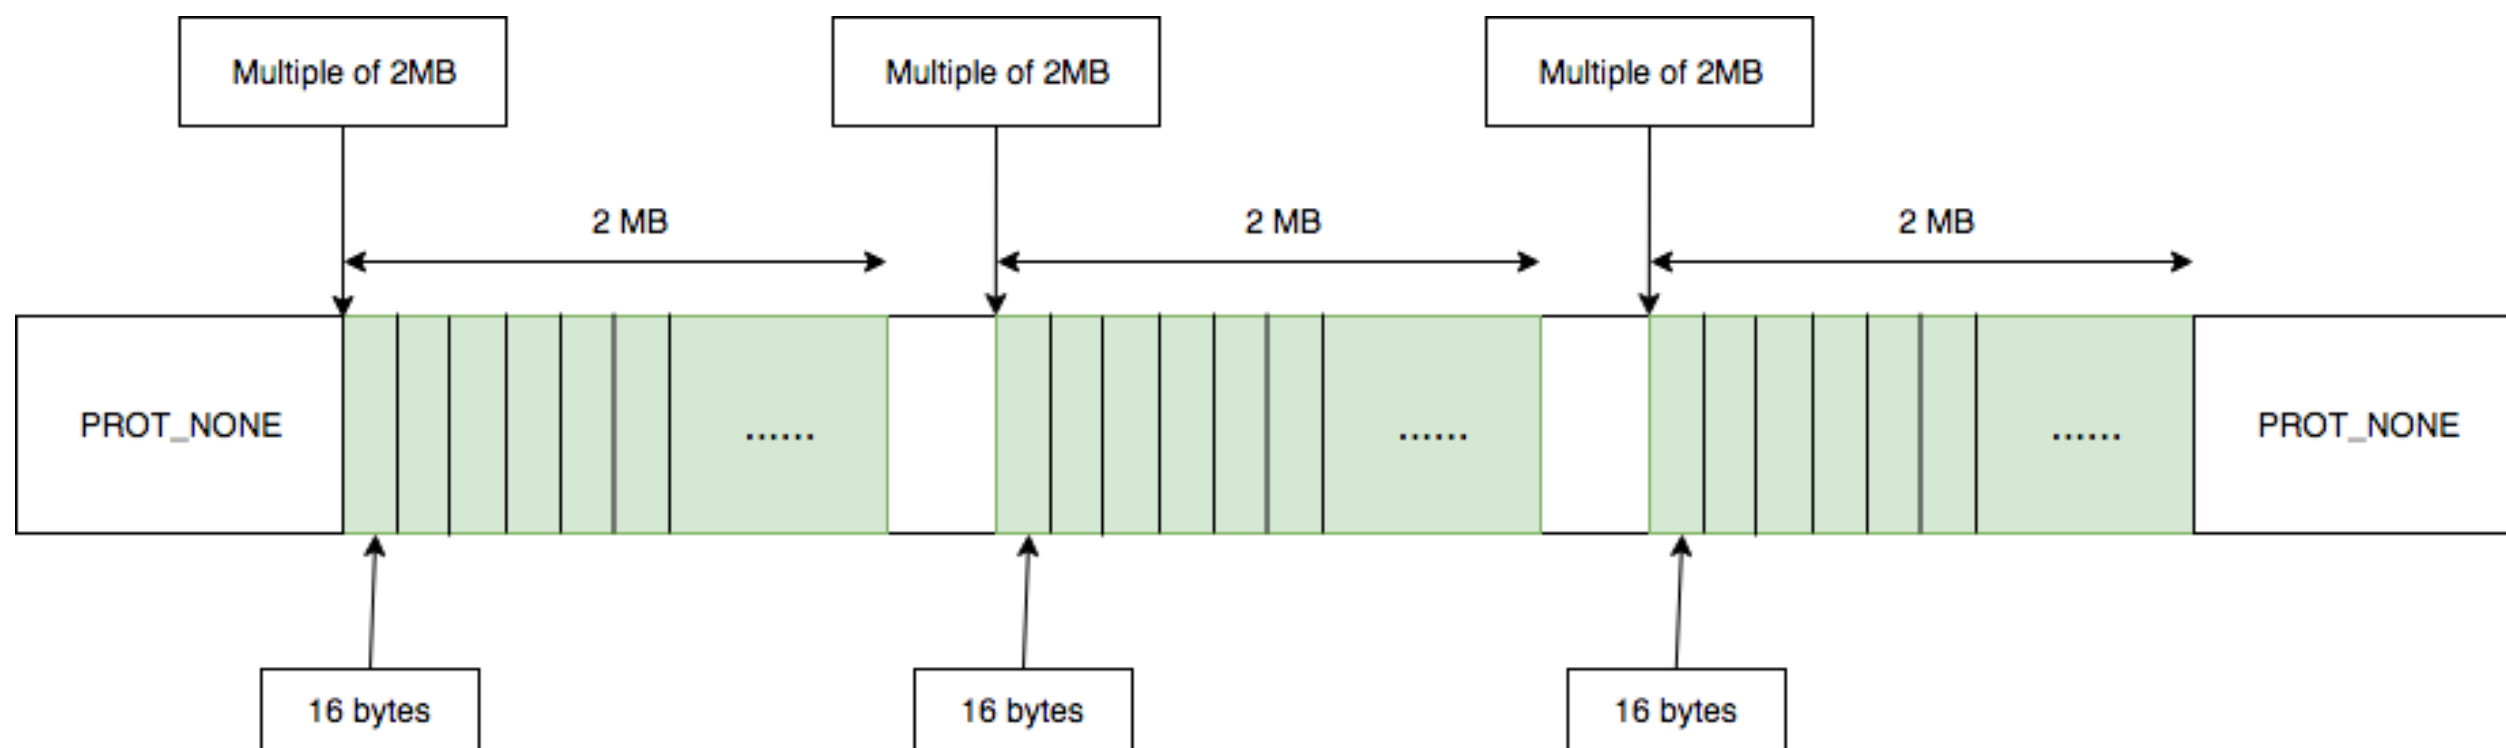
Allocation

- Return the first unused 16 byte block
- Remember, do not reuse memory

Deallocation

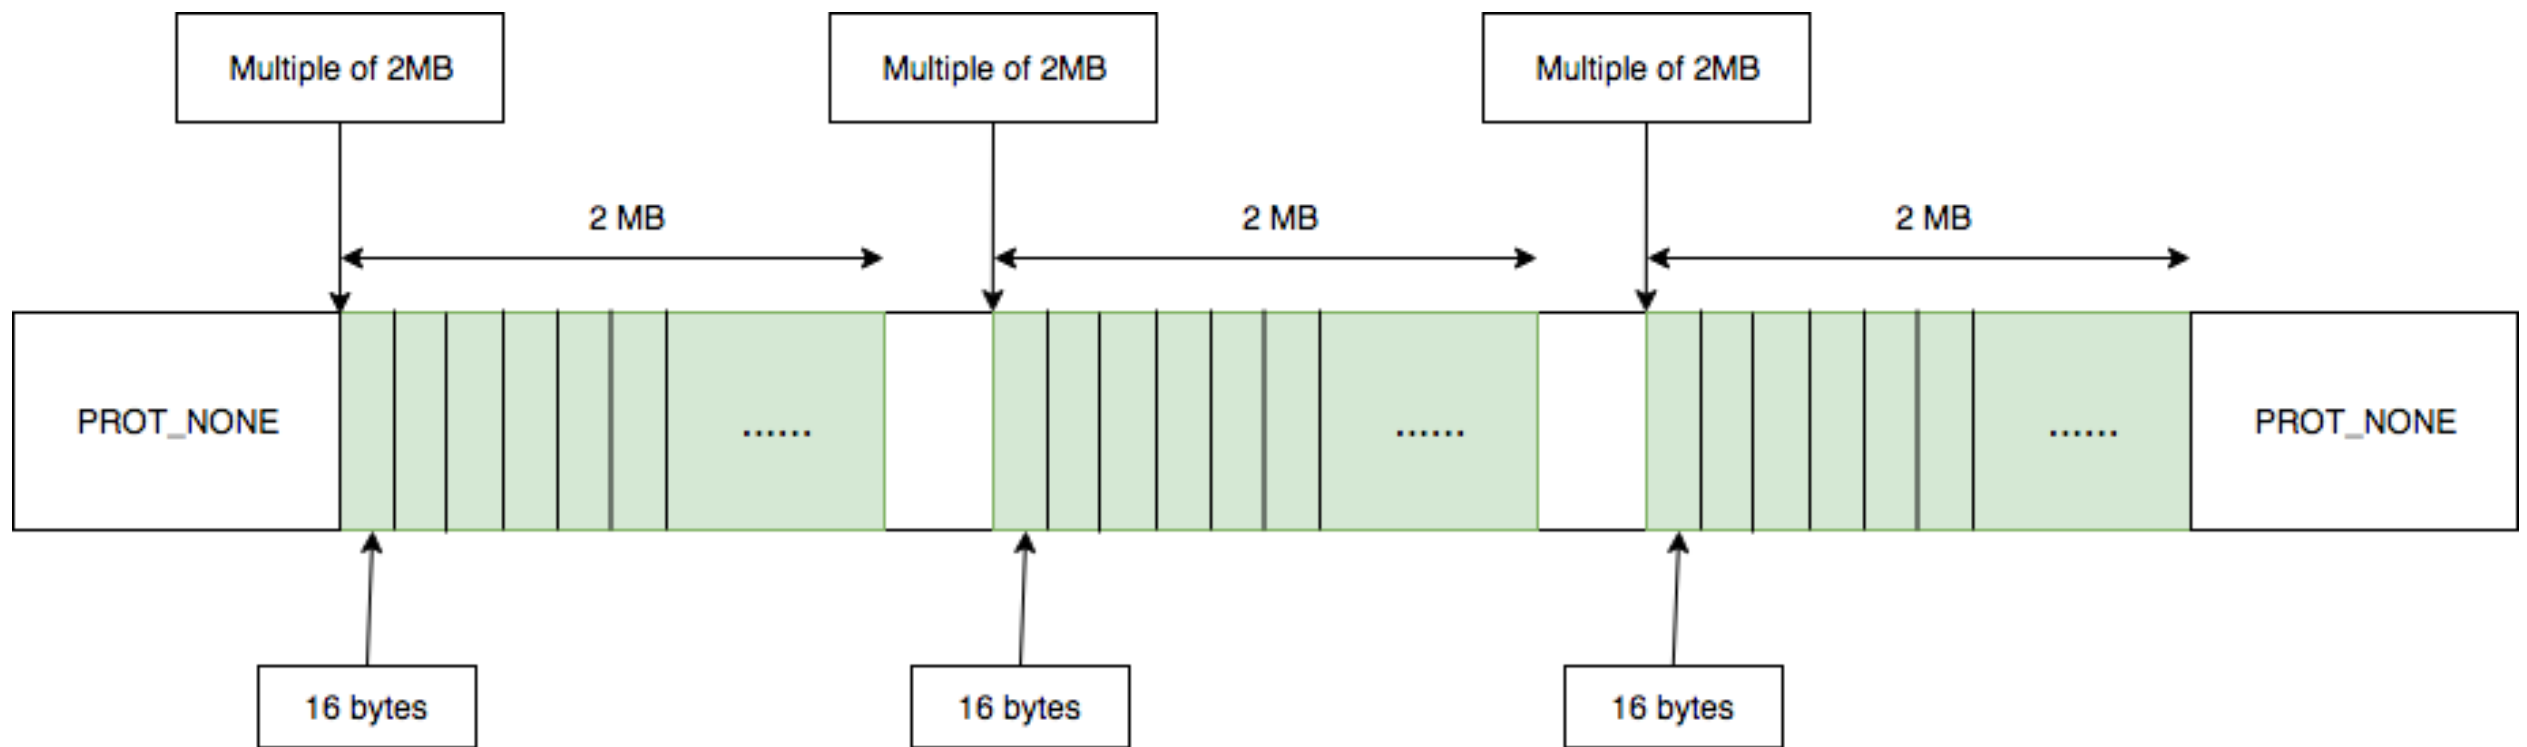
- Given a 16 byte block, how to quickly find the corresponding 2MB chunk?





Aligned Block List

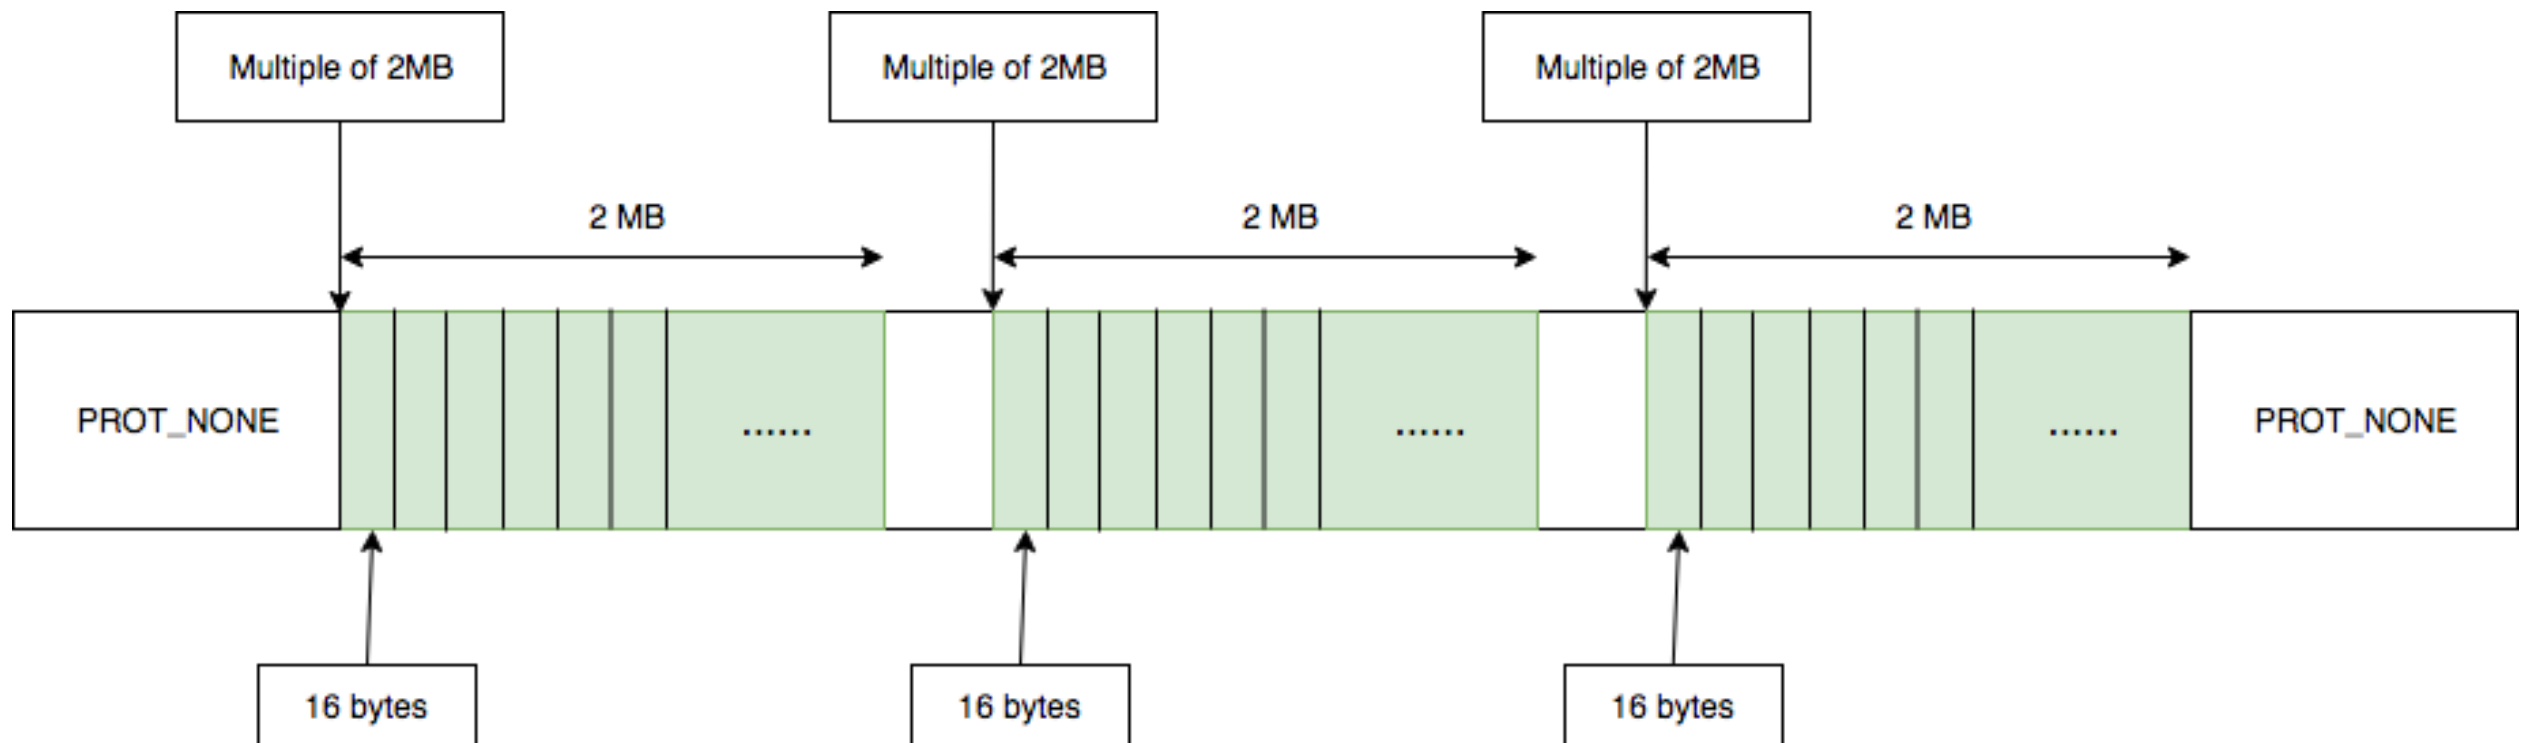
- Linked list of huge blocks, which are managed by another allocator (eg. BitmappedBlock)
- Length of each block is equal to its alignment, allowing for fast deallocations



- SafeBitmappedBlock
- AlignedBlockList
- AscendingPageAllocator

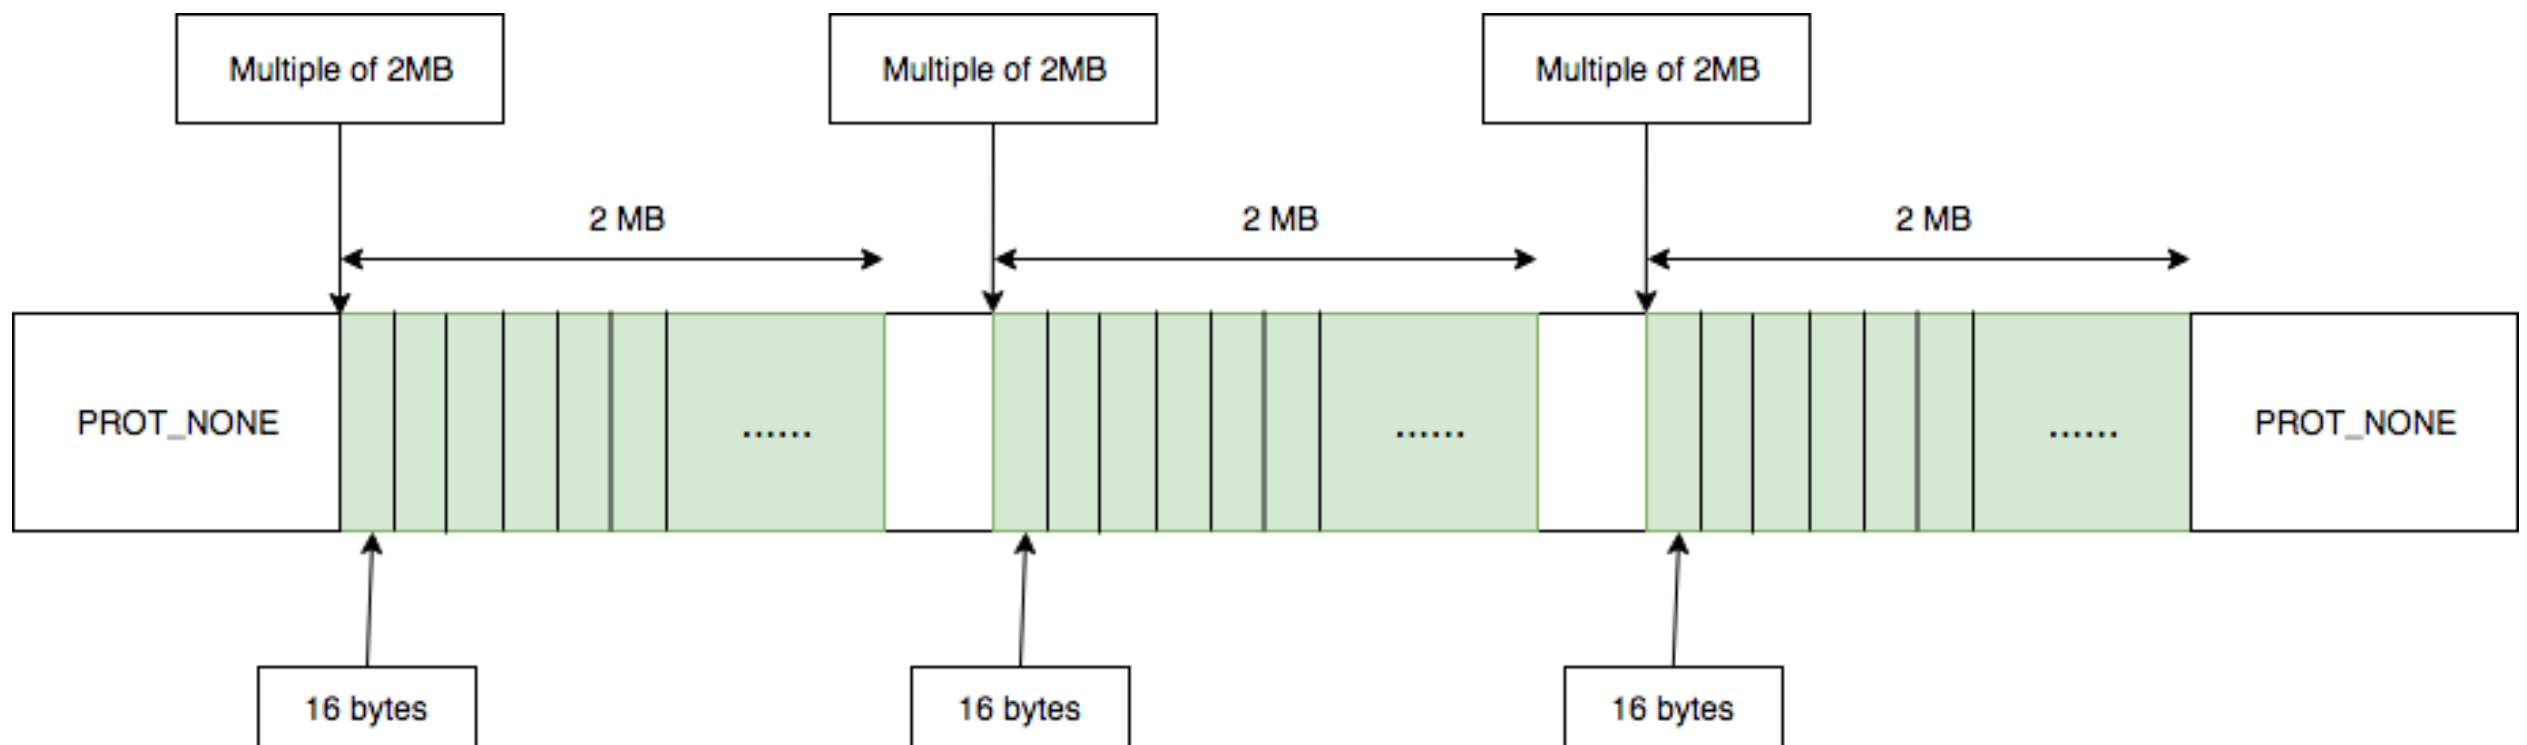
SafeAllocator v1.0

```
1 alias SafeAllocator = Segregator!(  
2     16, AlignedBlockList!(SafeBitmappedBlock!16, AscendingPageAllocator*, 1 << 21),  
3     AscendingPageAllocator*  
4 )
```



SafeAllocator v1.0

```
1 alias SafeAllocator = Segregator!(  
2     16, AlignedBlockList!(SafeBitmappedBlock!16, AscendingPageAllocator*, 1 << 21),  
3     32, AlignedBlockList!(SafeBitmappedBlock!32, AscendingPageAllocator*, 1 << 21),  
4     64, AlignedBlockList!(SafeBitmappedBlock!64, AscendingPageAllocator*, 1 << 21),  
5     128, AlignedBlockList!(SafeBitmappedBlock!128, AscendingPageAllocator*, 1 << 21),  
6     ...  
7     AscendingPageAllocator*  
8 )
```



```
1 SafeAllocator safeAlloc;
2 // initialization of safeAlloc
3 // ...
4
5 Point* p = safeAlloc.make!Point();
6 // do stuff with point
7 safeAlloc.dispose(p);
8
9 User* u = safeAlloc.make!User();
10 // ...
11
12 p.x = 100; // either crashes or modifies a valid Point object
13 u.name[0] = 'a'; // changes the first letter to 'a'
```

SafeAllocator v1.0

```
1 alias SafeAllocator = Segregator!(  
2     16, AlignedBlockList!(SafeBitmappedBlock!16, AscendingPageAllocator*, 1 << 21),  
3     32, AlignedBlockList!(SafeBitmappedBlock!32, AscendingPageAllocator*, 1 << 21),  
4     64, AlignedBlockList!(SafeBitmappedBlock!64, AscendingPageAllocator*, 1 << 21),  
5     128, AlignedBlockList!(SafeBitmappedBlock!128, AscendingPageAllocator*, 1 << 21),  
6     ...  
7     AscendingPageAllocator*  
8 )
```

- Safe (does not reuse memory)
- High fragmentation
- Lots of page faults

SafeAllocator v1.0

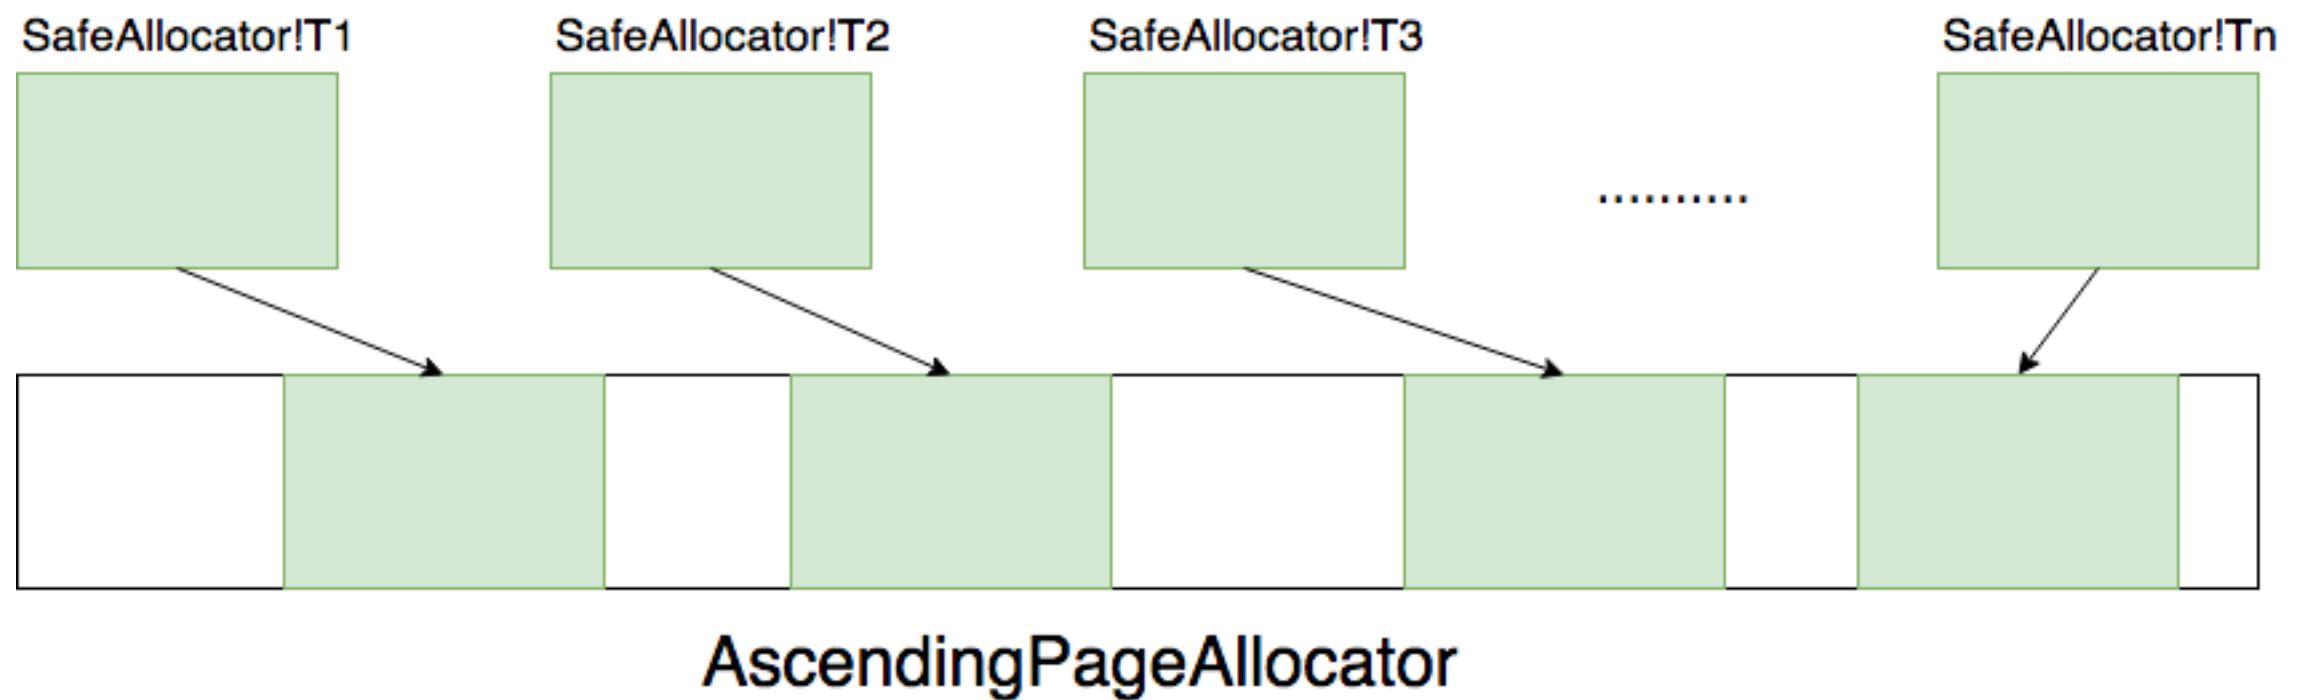
```
1 alias SafeAllocator = Segregator!(  
2     16, AlignedBlockList!(SafeBitmappedBlock!16, AscendingPageAllocator*, 1 << 21),  
3     32, AlignedBlockList!(SafeBitmappedBlock!32, AscendingPageAllocator*, 1 << 21),  
4     64, AlignedBlockList!(SafeBitmappedBlock!64, AscendingPageAllocator*, 1 << 21),  
5     128, AlignedBlockList!(SafeBitmappedBlock!128, AscendingPageAllocator*, 1 << 21),  
6     ...  
7     AscendingPageAllocator*  
8 )
```

- Safe (does not reuse memory)
- High fragmentation
- Lots of page faults

Not reusing addresses is safe but inefficient

```
1 struct Point
2 {
3     int x;
4     int y;
5 }
6
7 struct User
8 {
9     char[] name;
10    int age;
11 }
12
13 Point* p = Mallocator.instance.make!Point();
14 // do stuff with point
15 Mallocator.instance.dispose(p);
16
17 User* u = Mallocator.instance.make!User();
18 // ...
19
20 p.x = 100;
21 u.name[0] = 'a';
22
```


Reuse memory only for the same types



```
1 struct SafeAllocator(T)
2 {
3     alias blockSize = roundUpToPowerOf2(stateSize!T);
4
5     alias AllocType = AlignedBlockList!(
6         BitmappedBlock!blockSize,
7         AscendingPageAllocator*,
8         1 << 21
9     );
10
11 // ...
12 }
```

```

1 SafeAllocator!T* getSafeAllocator(T)()
2 {
3     static SafeAllocator!T safeAllocator;
4
5     if (!safeAllocator.isInit())
6         safeAllocator.initialize();
7
8     return &safeAllocator;
9 }
10

```

```

1 auto safePointAlloc = getSafeAllocator!Point();
2 auto safeUserAlloc = getSafeAllocator!User();
3
4 Point* p = safePointAlloc.make!Point();
5 // do stuff with p
6 safePointAlloc.dispose(p);
7
8 User* u = safeUserAlloc.make!User();
9 // ...
10 p.x = 100; // this is now safe

```

```
1 struct SafeAllocator(T)
2 {
3     alias blockSize = roundUpToPowerOf2(stateSize!T);
4
5     alias AllocType = AlignedBlockList!(
6         BitmappedBlock!blockSize,
7         SharedAscendingPageAllocator*,
8         1 << 21
9     );
10
11 // ...
12 }
```

SafeAllocator v2.0

Safe

High performance

Scalable

SafeAllocator v2.0

Safe

High performance

Scalable

Many types = high fragmentation



```
1 struct Point
2 {
3     int x;
4     int y;
5 }
6
7
8 struct Interval
9 {
10    int left;
11    int right;
12 }
13
```


Memory layout

Many types = High fragmentation

Reuse memory among types with the same memory layout

Layout

```
1 struct Point
2 {
3     int x;
4     int y;
5 }
6
7 Layout!Point = [0, int, 4, int, 8]
8
9 struct User
10 {
11     char[] name;
12     int age;
13 }
14
15 Layout!User = [0, char[], 16, int, 24]
16
```

Layout

```
1 struct User
2 {
3     char[] name;
4     int age;
5 }
6
7 Layout!User = [0, char[], 16, int, 24];
8
9 struct Stock
10 {
11     char[] id;
12     int price;
13 }
14
15 Layout!Stock = [0, char[], 16, int, 24];
16
```

Layout


```
1 align(16) struct User
2 {
3     char[] name;
4     int age;
5 }
6
7 Layout!User = [0, char[], 16, int, 32];
8
9 struct Stock
10 {
11     char[] id;
12     int price;
13 }
14
15 Layout!Stock = [0, char[], 16, int, 24];
16
```

SafeAllocator v3.0


```
1 struct SafeAllocator(T...)
2 {
3     alias blockSize = roundUpToPowerOf2(stateSize!T[$ - 1]);
4
5     alias AllocType = AlignedBlockList!(
6         BitmappedBlock!blockSize,
7         SharedAscendingPageAllocator*,
8         1 << 21
9     );
10
11 // ...
12 }
```

```
1 SafeAllocator!T* getSafeAllocator(T)()
2 {
3     static SafeAllocator!(Layout!T) safeAllocator;
4
5     if (!safeAllocator.isInit())
6         safeAllocator.initialize();
7
8     return &safeAllocator;
9 }
```


16 byte allocations




16 byte allocations



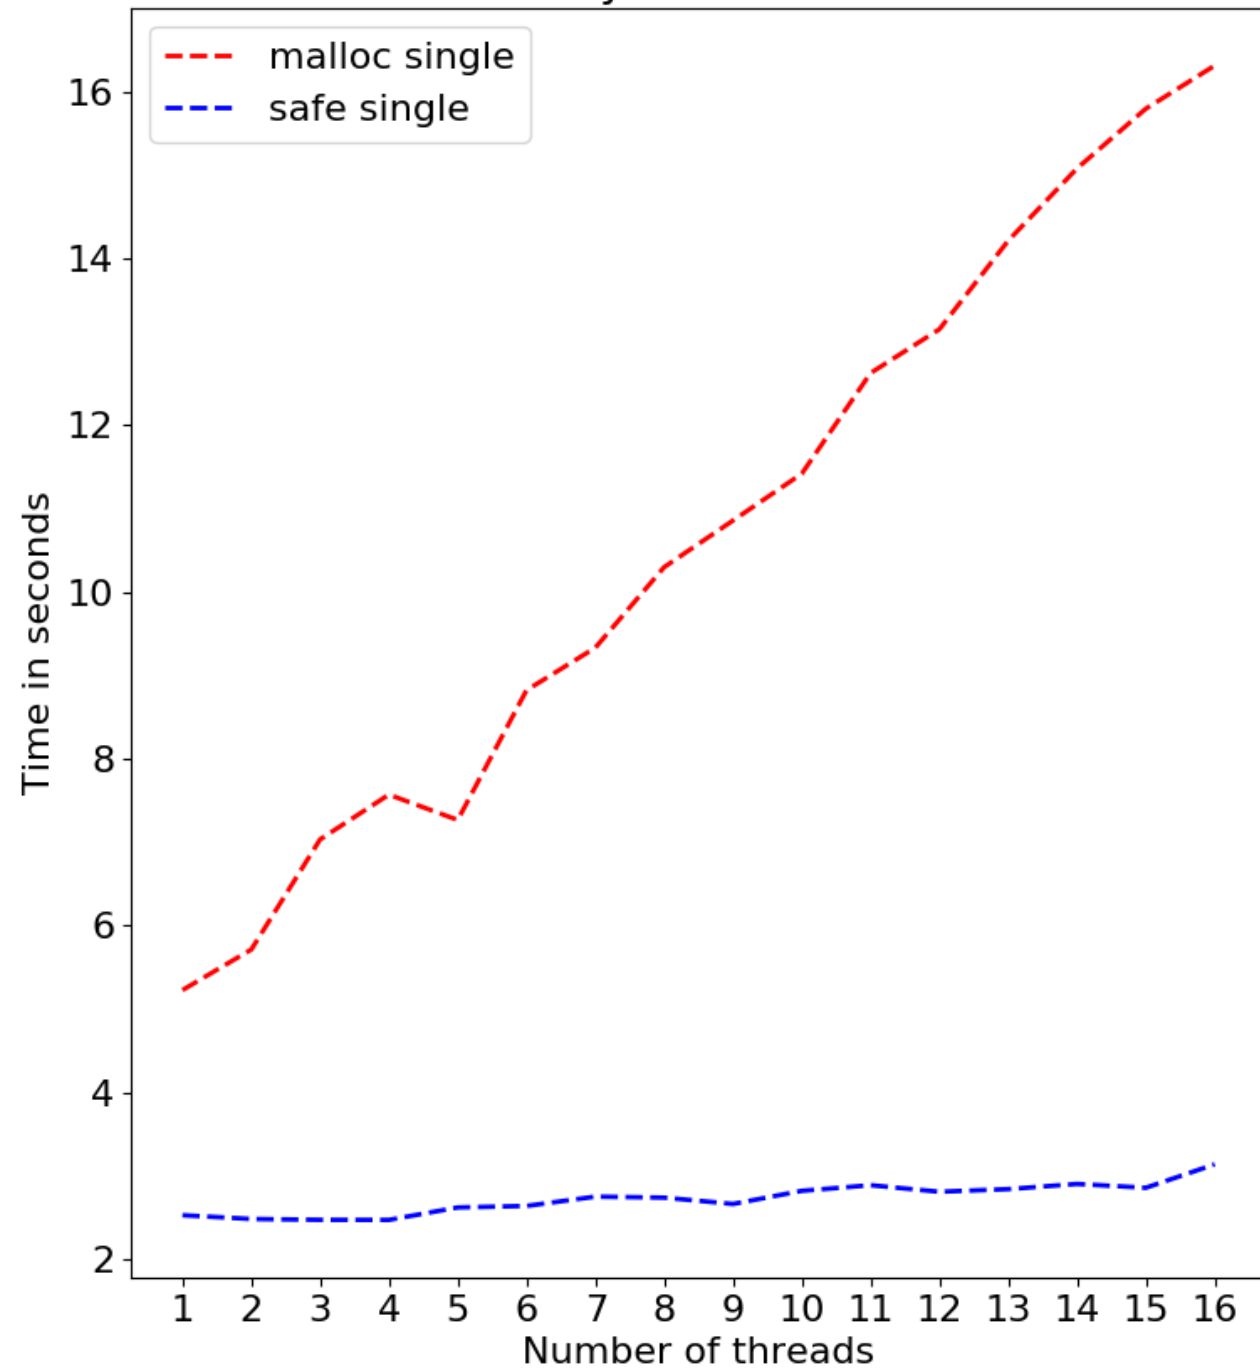
64 byte allocations




64 byte allocations



512 byte allocations



512 byte allocations



Summary

	SAFETY	PERFORMANCE	FRAGMENTATION
SAFE ALLOCATOR V1.0	Not reuse memory	Fast for small and large allocations	High
SAFE ALLOCATOR V2.0	Type reuse	Fast	Moderate
SAFE ALLOCATOR V3.0	Layout reuse	Fast	Low

Future work

Improve documentation for current allocators

Improve specification for Layout and SafeAllocator

More benchmarks

Acknowledgments