

Compacting the Uncompactable!

MESH Compacting Memory Management
For C/C++ Applications

Emery Berger
UMass Amherst / MSR

with Bobby Powers, David Tench, & Andrew McGregor
University of Massachusetts Amherst

<http://libmesh.org>

[PLDI 2019]



PIANO - VOCAL - GUITAR



The Impossible Dream (The Quest)

Lyrics by JOE DARIEN
Music by MITCH LEIGH



PIANO - VOCAL - GUITAR



The Impossible Dream (The Quest)

Lyrics by JOE DARIEN
Music by MITCH LEIGH



Reconquer all of Spain!



PIANO - VOCAL - GUITAR



The Impossible Dream (The Quest)

Lyrics by JOE DARIEN
Music by MITCH LEIGH

MALLOCC

OF
LA

MANGA



Reconquer all of Spain!



PIANO - VOCAL - GUITAR



The Impossible Dream (The Quest)

Lyrics by JOE DARIEN
Music by MITCH LEIGH

MALLOC

OF
LA

MANGA



~~Reconquer Malloc~~ all of Spain!



PIANO - VOCAL - GUITAR



The Impossible Dream (The Quest)

Lyrics by JOE DARIEN
Music by MITCH LEIGH

MALLOCC

OF
LA

MANGA



PIANO - VOCAL - GUITAR



The Impossible Dream (The Quest)

Lyrics by JOE DARIEN
Music by MITCH LEIGH

MALLOC

OF
LA

MANGA



(malloc)



PIANO - VOCAL - GUITAR



The Impossible Dream (The Quest)

Lyrics by JOE DARIEN
Music by MITCH LEIGH

MALLOCC

OF
LA
MANGA



PIANO - VOCAL - GUITAR



The Impossible Dream (The Quest)

Lyrics by JOE DARIEN
Music by MITCH LEIGH

MALLOCC

OF
LA
MANGA



PIANO - VOCAL - GUITAR



The Impossible Dream (The Quest)

Lyrics by JOE DARIEN
Music by MITCH LEIGH

MALLOCC

OF
LA
MANGA



PIANO - VOCAL - GUITAR



The Impossible Dream (The Quest)

Lyrics by JOE DARIEN
Music by MITCH LEIGH

MALLOCC

OF
LA

MANGA



PIANO - VOCAL - GUITAR



The Impossible Dream (The Quest)

Lyrics by JOE DARIEN
Music by MITCH LEIGH

MALLOCC

OF
LA

MANGA



PIANO - VOCAL - GUITAR



The Impossible Dream (The Quest)

Lyrics by JOE DARIEN
Music by MITCH LEIGH

MALLOCC

OF
LA
MANGA



?



PIANO - VOCAL - GUITAR



The Impossible Dream (The Quest)

Lyrics by JOE DARIEN
Music by MITCH LEIGH

MALLOCC

OF
LA
MANGA



?



PIANO - VOCAL - GUITAR



The Impossible Dream (The Quest)

Lyrics by JOE DARIEN
Music by MITCH LEIGH

MALLOCC

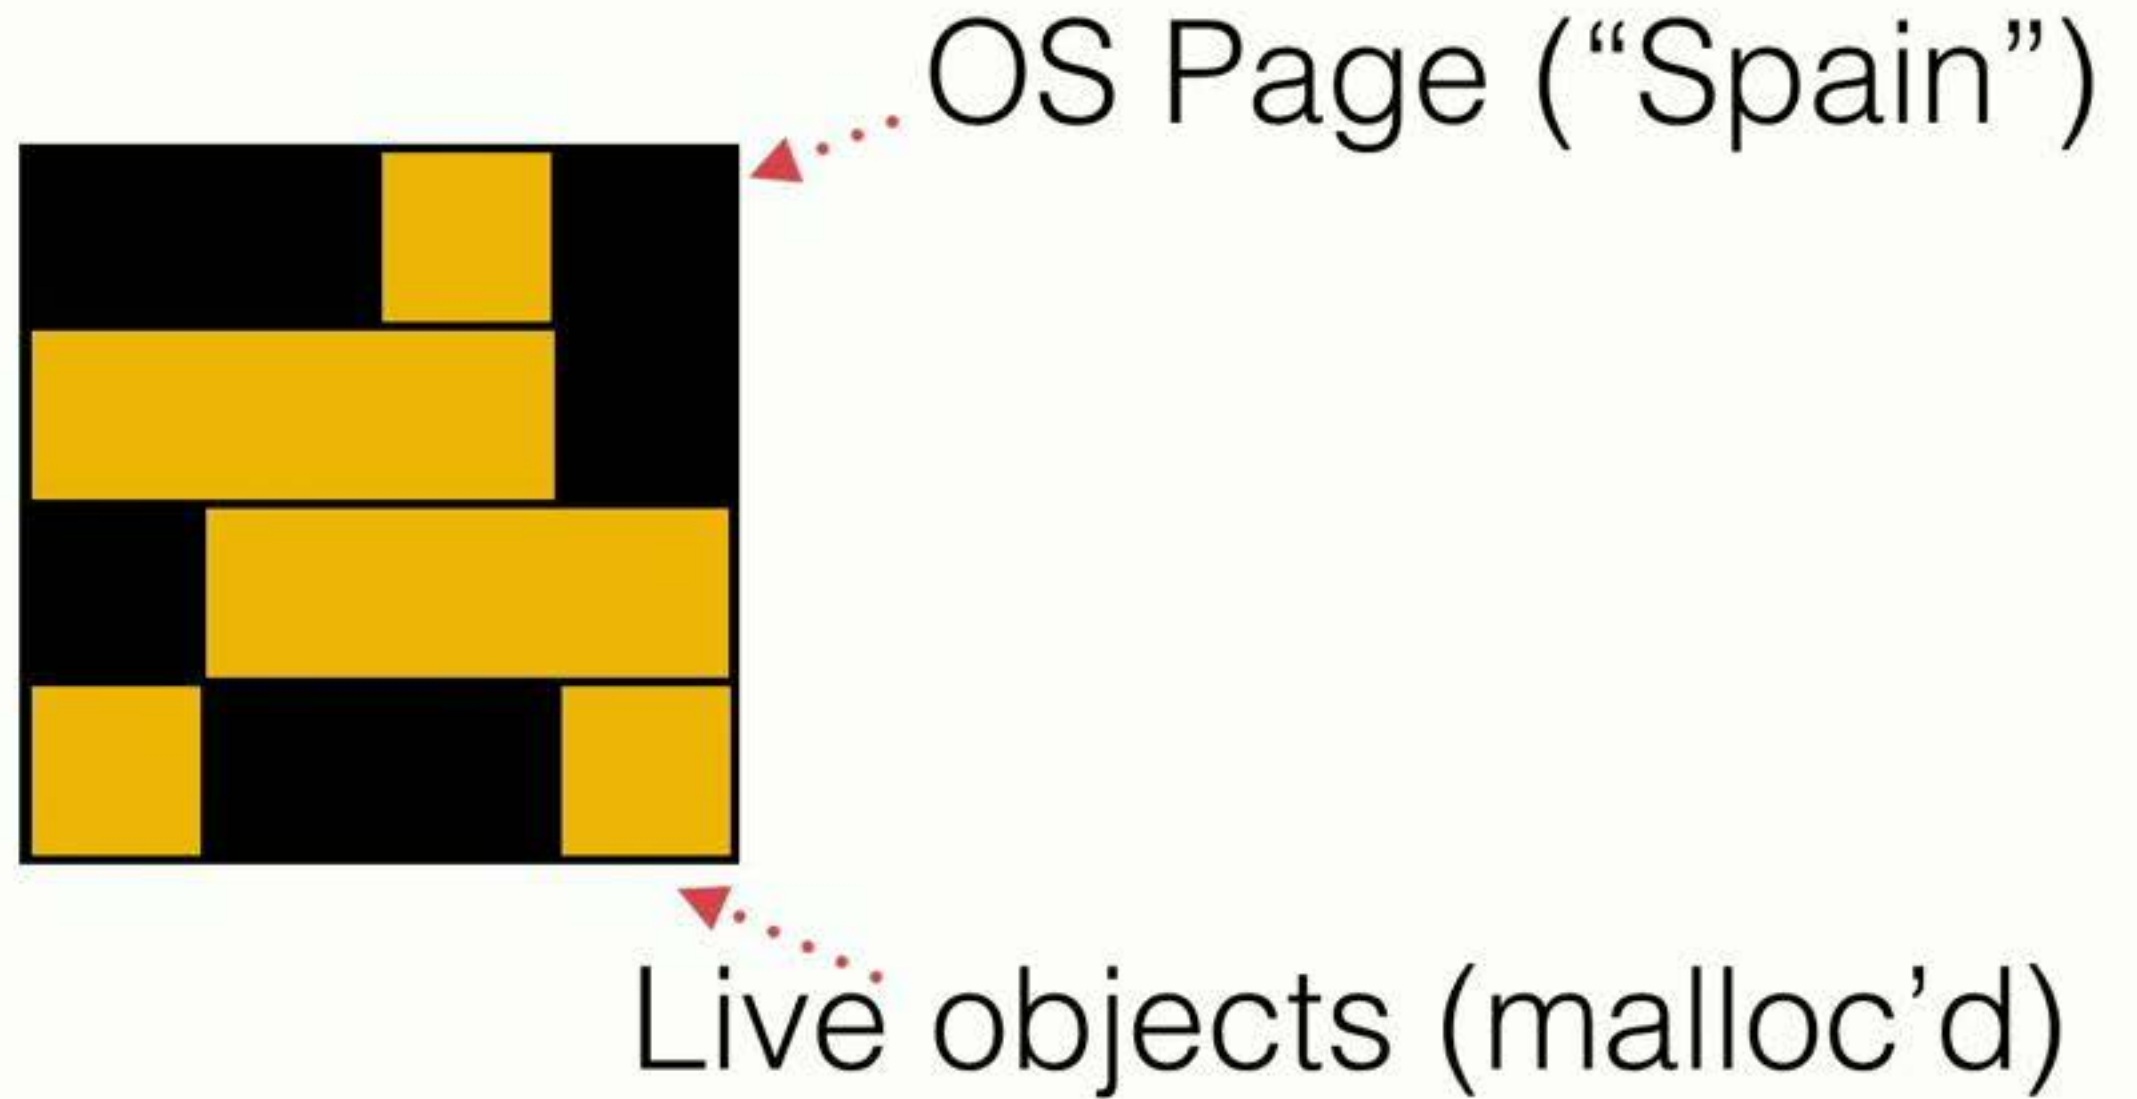
OF
LA
MANGA






i!

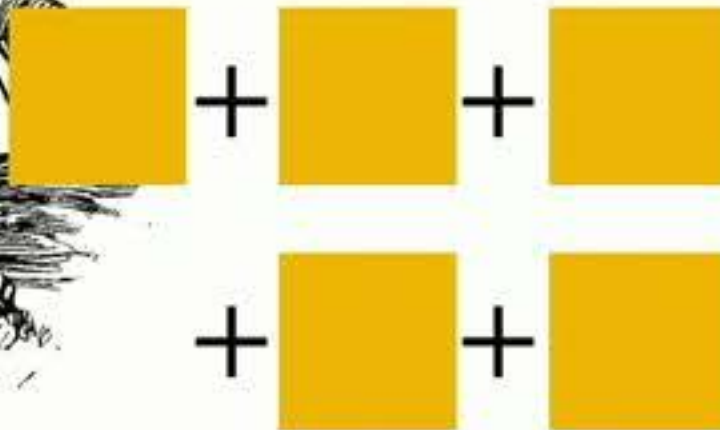

¡Fragmentación!



¡Fragmentación!




OS Pages





Live objects

¡Fragmentación!




OS Pages



Live objects

» 1

¡Fragmentación!





OS Pages


$O(\log \frac{\text{large yellow rectangle}}{\text{small yellow square}})$ 13x!

[Robson '77]

Live objects



¿Compacción?





OS Pages


$O(\log \frac{\text{large yellow rectangle}}{\text{small yellow square}})$ 13x!

[Robson '77]

Live objects





¿Compacción?



OS Pages


$O(1)$



Live objects



i!





NG
E

M. Ritchie




NG
E

M. Ritchie

THE
C
PROGRAMMING
LANGUAGE


Brian W. Kernighan • Dennis M. Ritchie

PRENTICE HALL SOFTWARE SERIES



The Swift
Programming
Language
Swift 5 Edition





[Lattner, 2016]

Why not a tracing GC?

- Native interoperability with unmanaged code
- Deterministic destruction provides:
 - No “finalizer problems” like resurrection, threading, etc.
 - Deterministic performance: can test/debug performance stutters
- Performance:
 - GC use ~3-4x more memory than ARC to achieve good performance
 - Memory usage is very important for mobile and cloud apps
 - Incremental/concurrent GCs slow the mutator like ARC does

Native interoperability with unmanaged code

Deterministic destruction provides:

- No “finalizer problems” like resurrection, threading, etc.
- Deterministic performance: can test/debug performance stutters

Performance:

- GC use ~3-4x more memory than ARC to achieve good performance
- Memory usage is very important for mobile and cloud apps
- Incremental/concurrent GCs slow the mutator like ARC does


Quantifying the Performance of Garbage Collection vs. Explicit Memory Management

Matthew Hertz, Emery D. Berger. OOPSLA'05

THE
C
PROGRAMMING
LANGUAGE




Brian W. Kernighan • Dennis M. Ritchie

PRENTICE HALL SOFTWARE SERIES






The Swift
Programming
Language
Swift 5 Edition









redis








redis

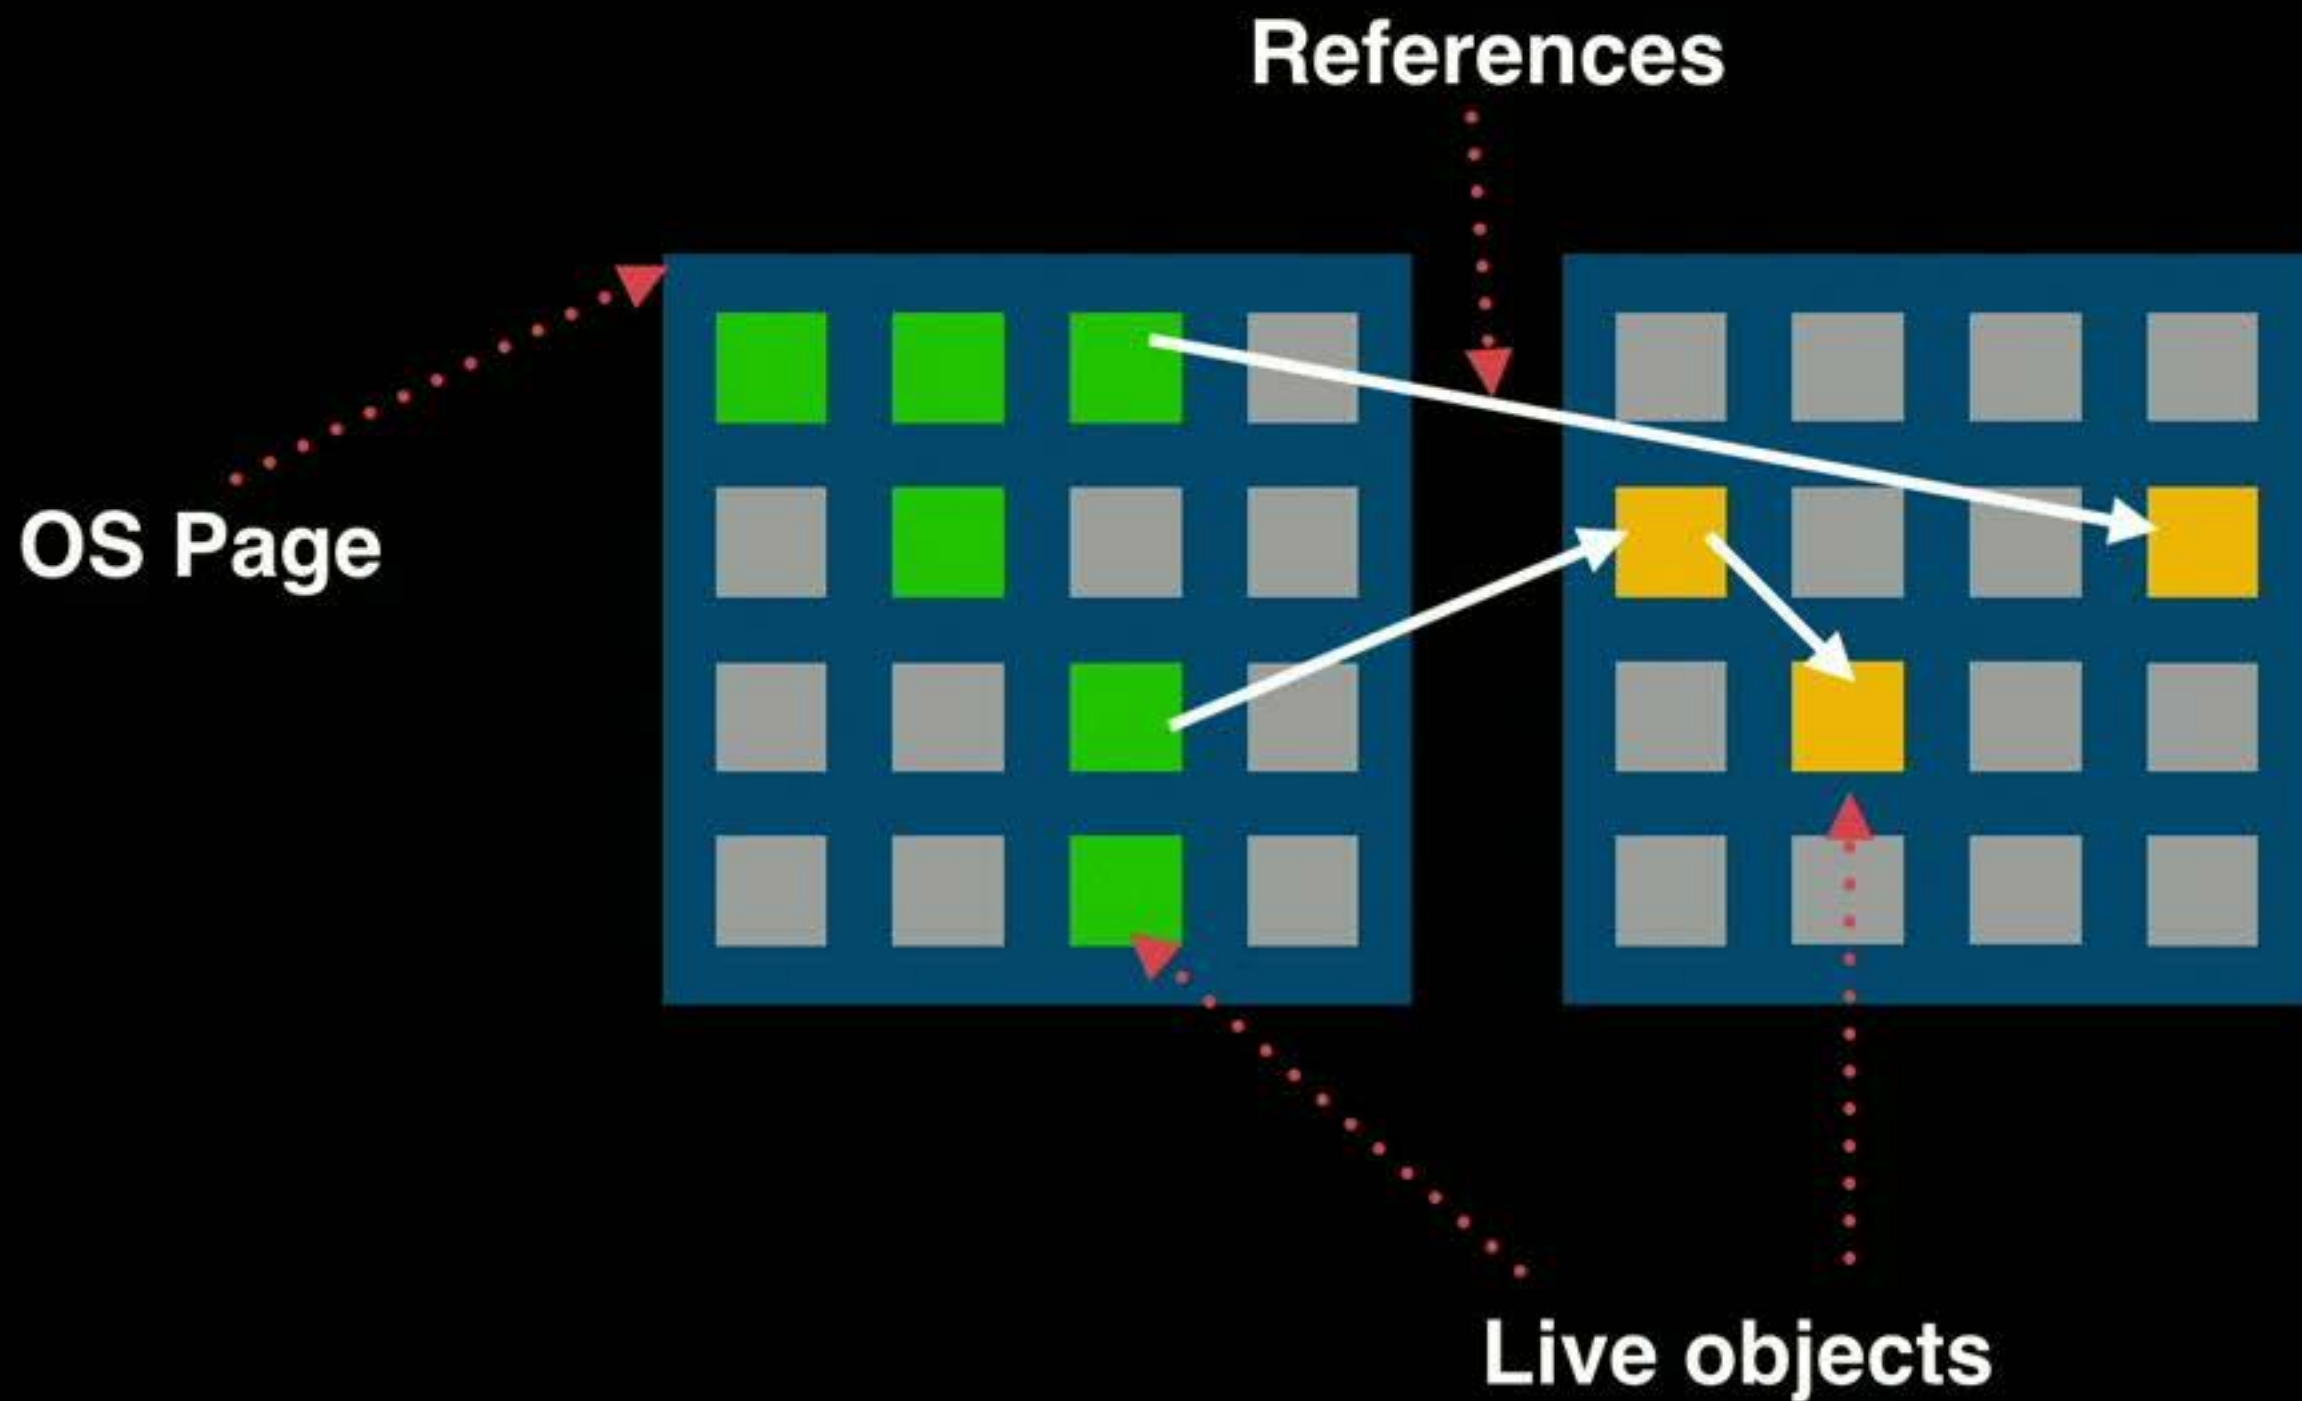


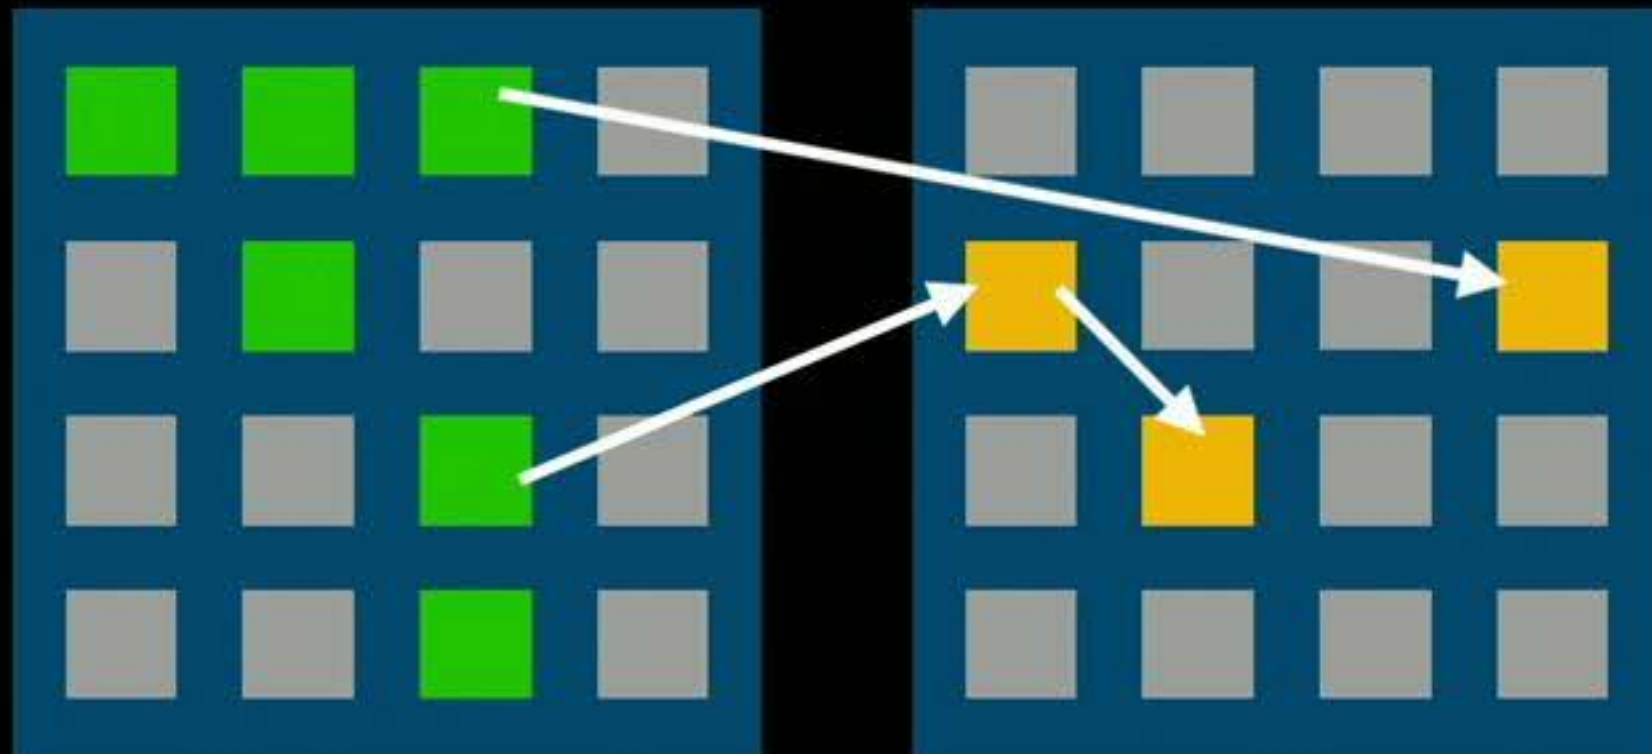




¡Compacción!


Compaction In Action






```
$ cc -o yolo main.cc
```

```
$ strip yolo
```



```
$ cc -o yolo main.cc  
$ strip yolo  
$ ./yolo
```

0xDEADC000



**No way to precisely
distinguish pointers
from integers**

0xBEEFC000




```
$ cc -o yolo main.cc
$ strip yolo
$ ./yolo
```



No way to precisely distinguish pointers from integers



0xDEADC000



0xBEEFC000



```
union tiny
{
    int * ptr;
    uintptr_t flag;
};
tiny x;

// initialize
x.ptr = new int;

// set flag true
x.flag |= 1;
```


MESH

*Compaction without
Relocation for C/C++*


No code changes

No recompilation

`LD_PRELOAD` and go

17% heap size reduction

< 1% performance overhead

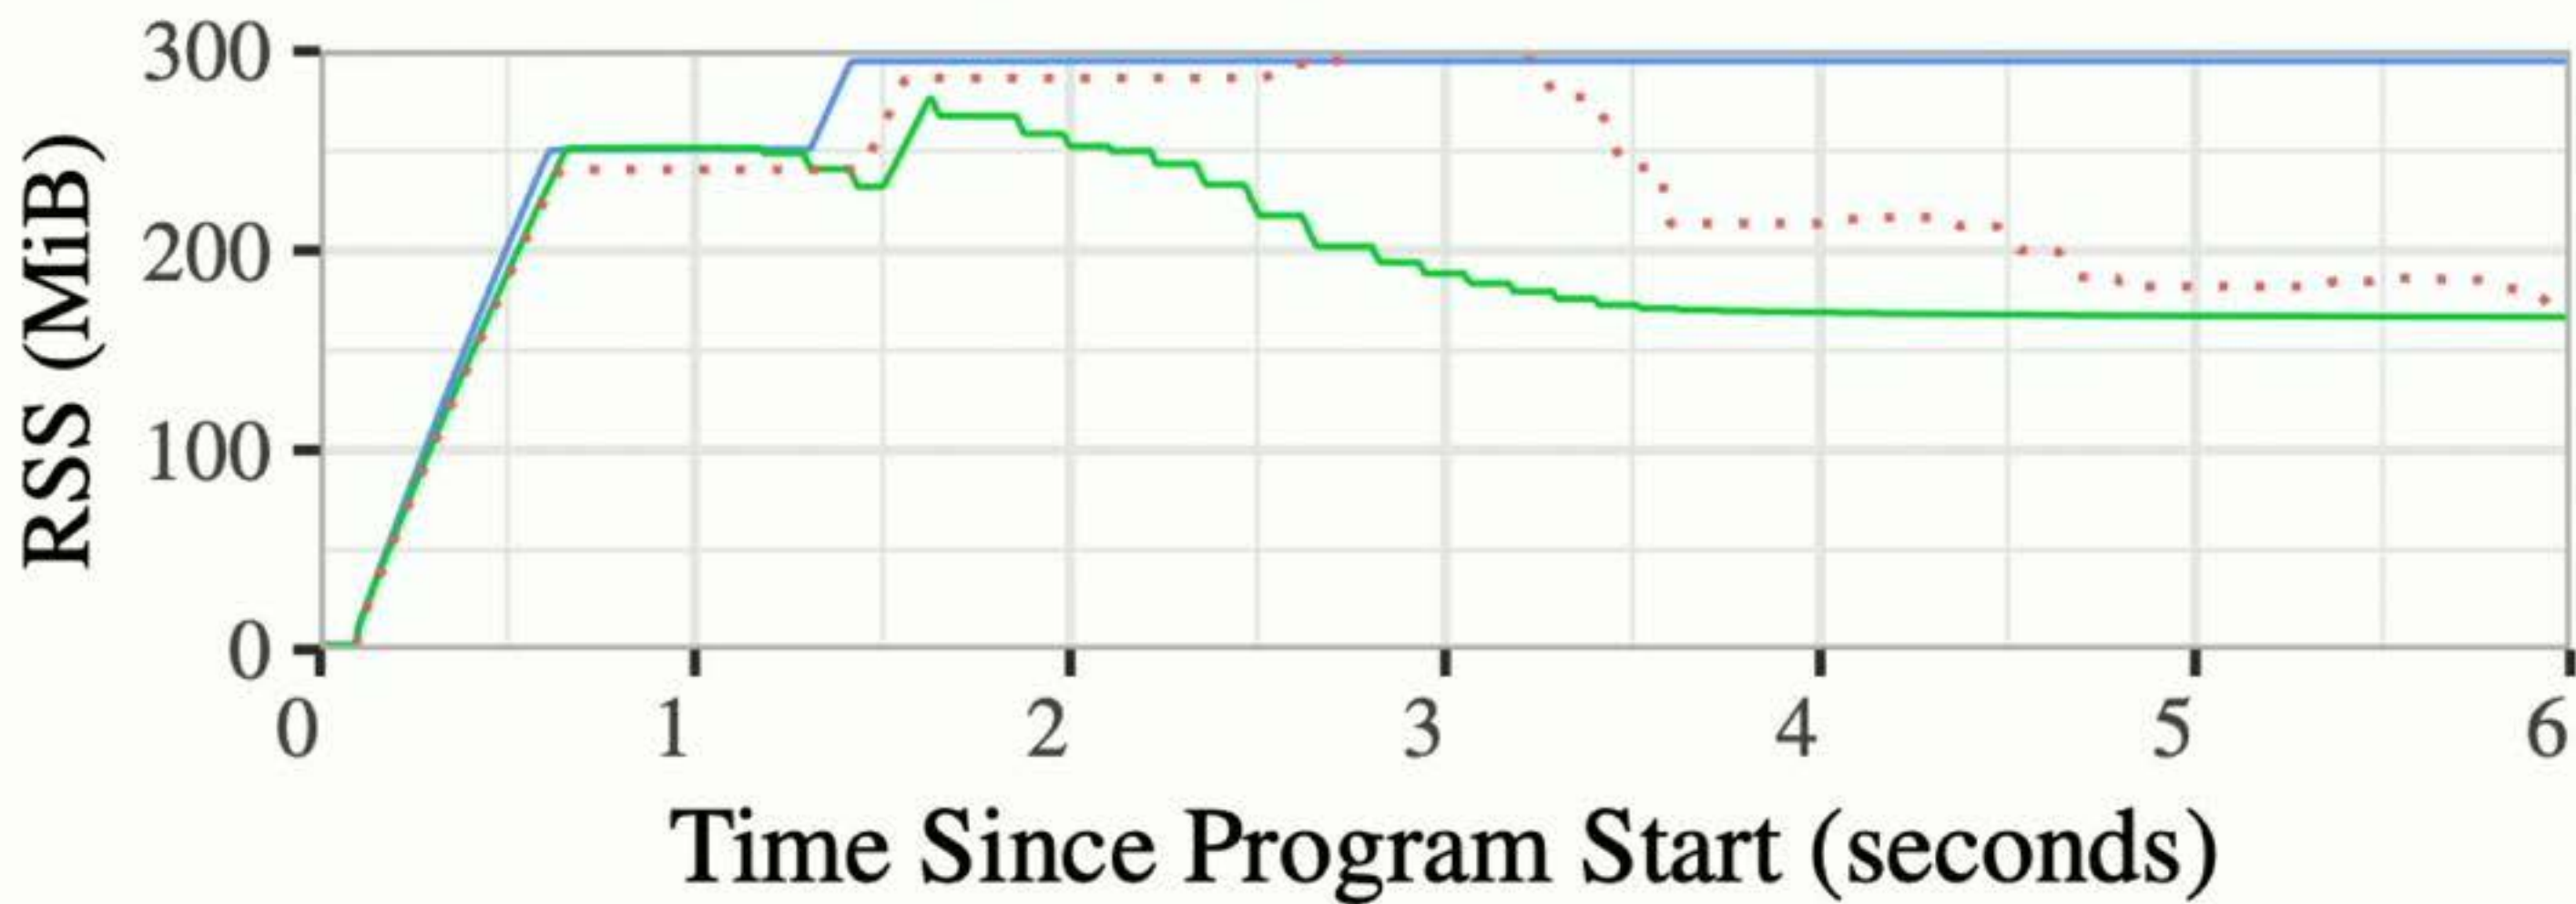


— default jemalloc — Mesh



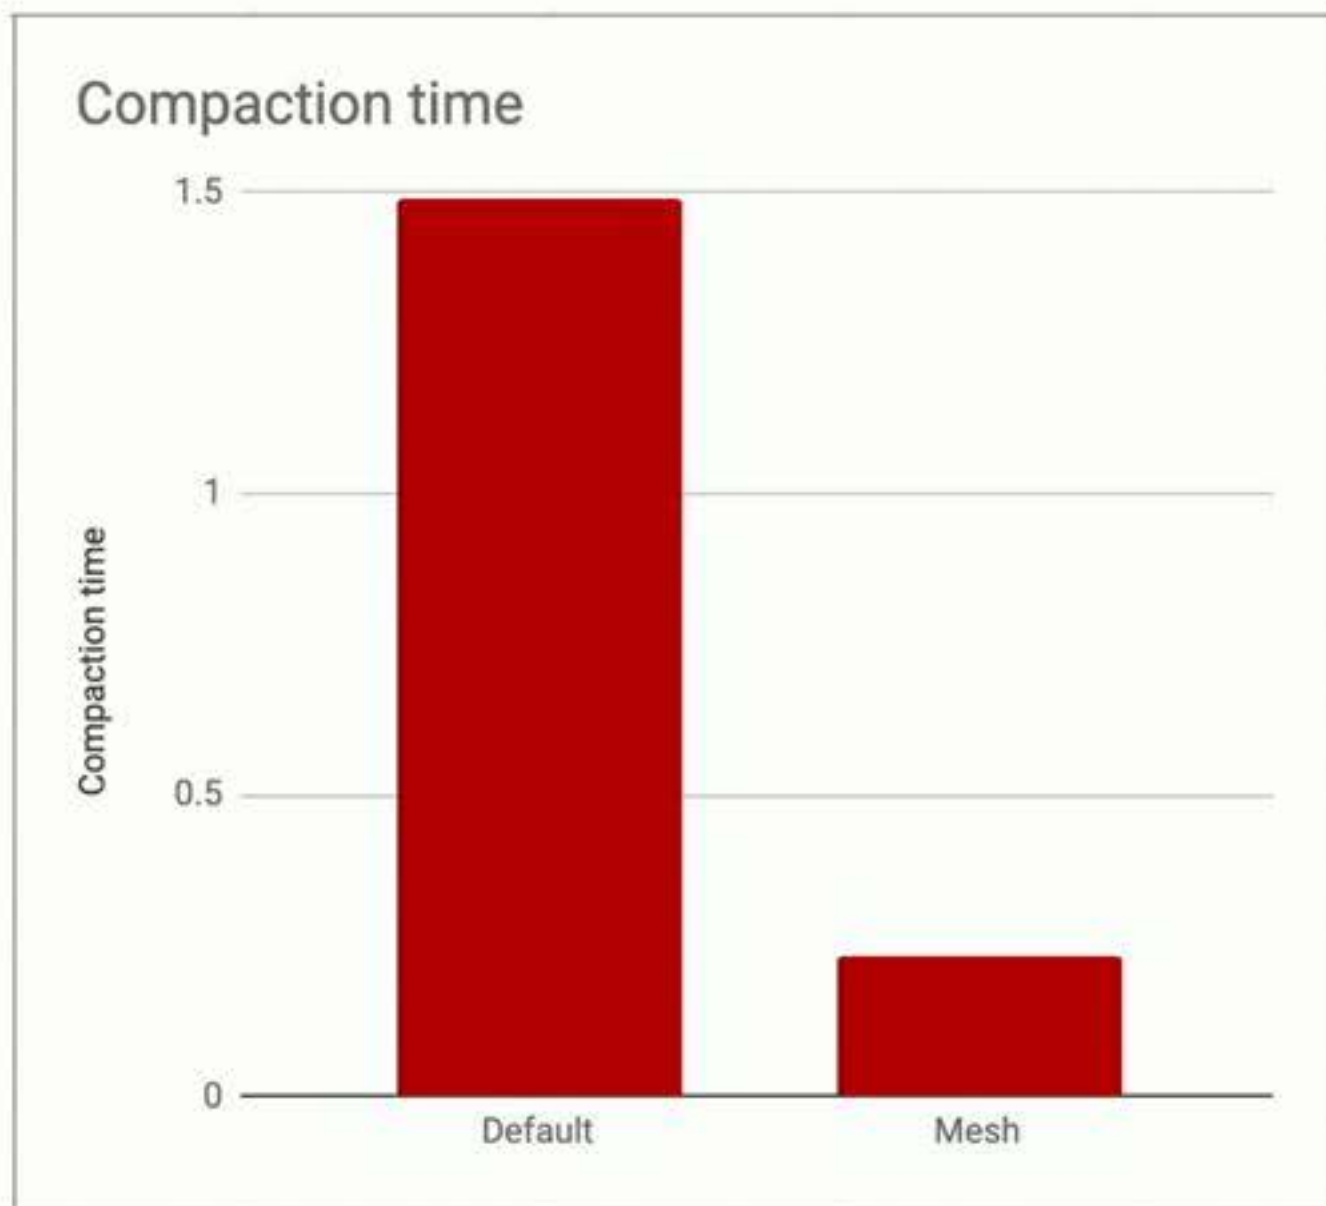
redis

- jemalloc + activedefrag
- Mesh
- no compaction






redis




5x reduction in
time spent
compacting

Meshing: *compaction without
(virtual!) relocation*



Pages are **Meshable** when they:


- Hold objects of the same size class






Pages are **Meshable** when they:

- Hold objects of the same size class
- Have non-overlapping object offsets

Meshing




Virtual




Physical

Meshing


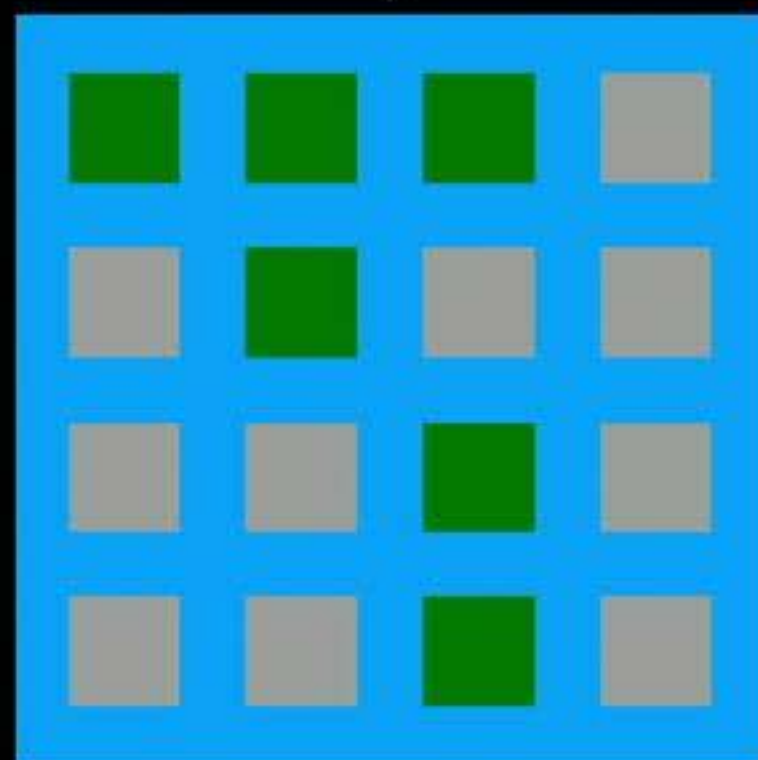


Mark virtual page read-only



Virtual




Physical





Meshing



Virtual



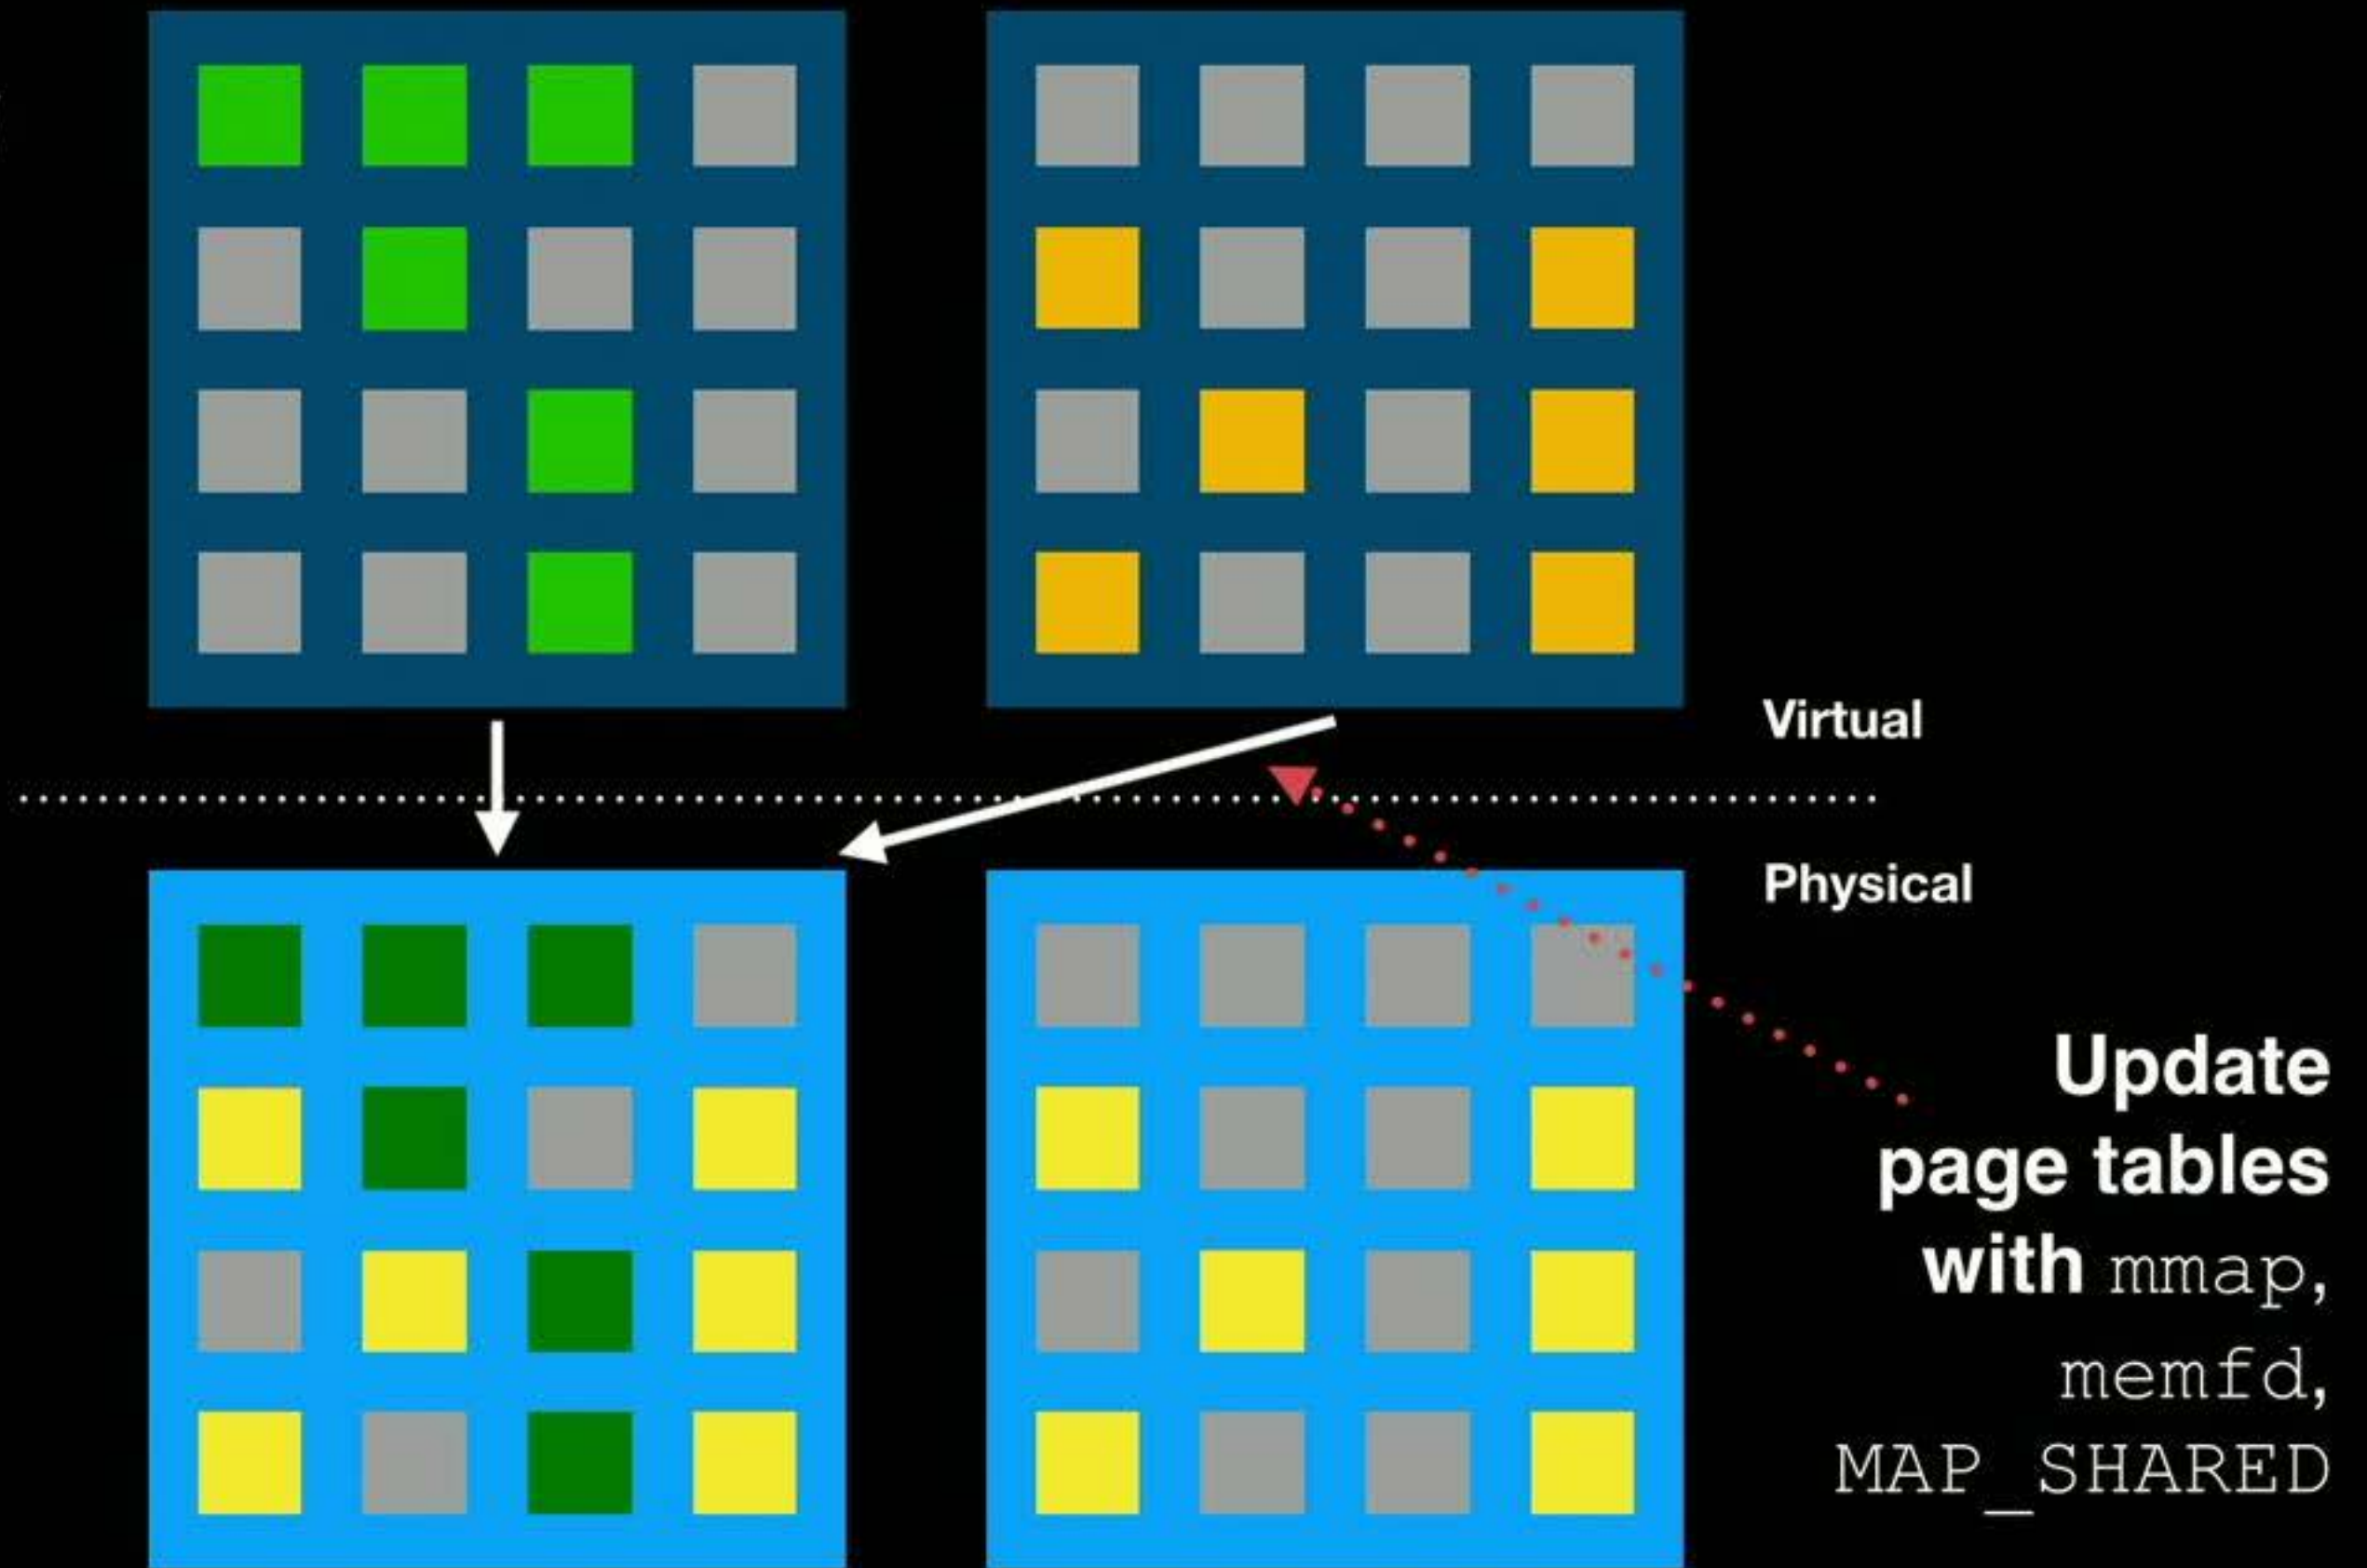
Physical





Copy
(maintaining
offsets)



Meshing



Meshing





Mark virtual page read/write


Virtual




Physical




Meshing



Virtual





Physical




Physical page returned to OS

Meshing




No virtual addresses changed!

Virtual





Physical



Worst Case:

low occupancy,
non-meshable
pages



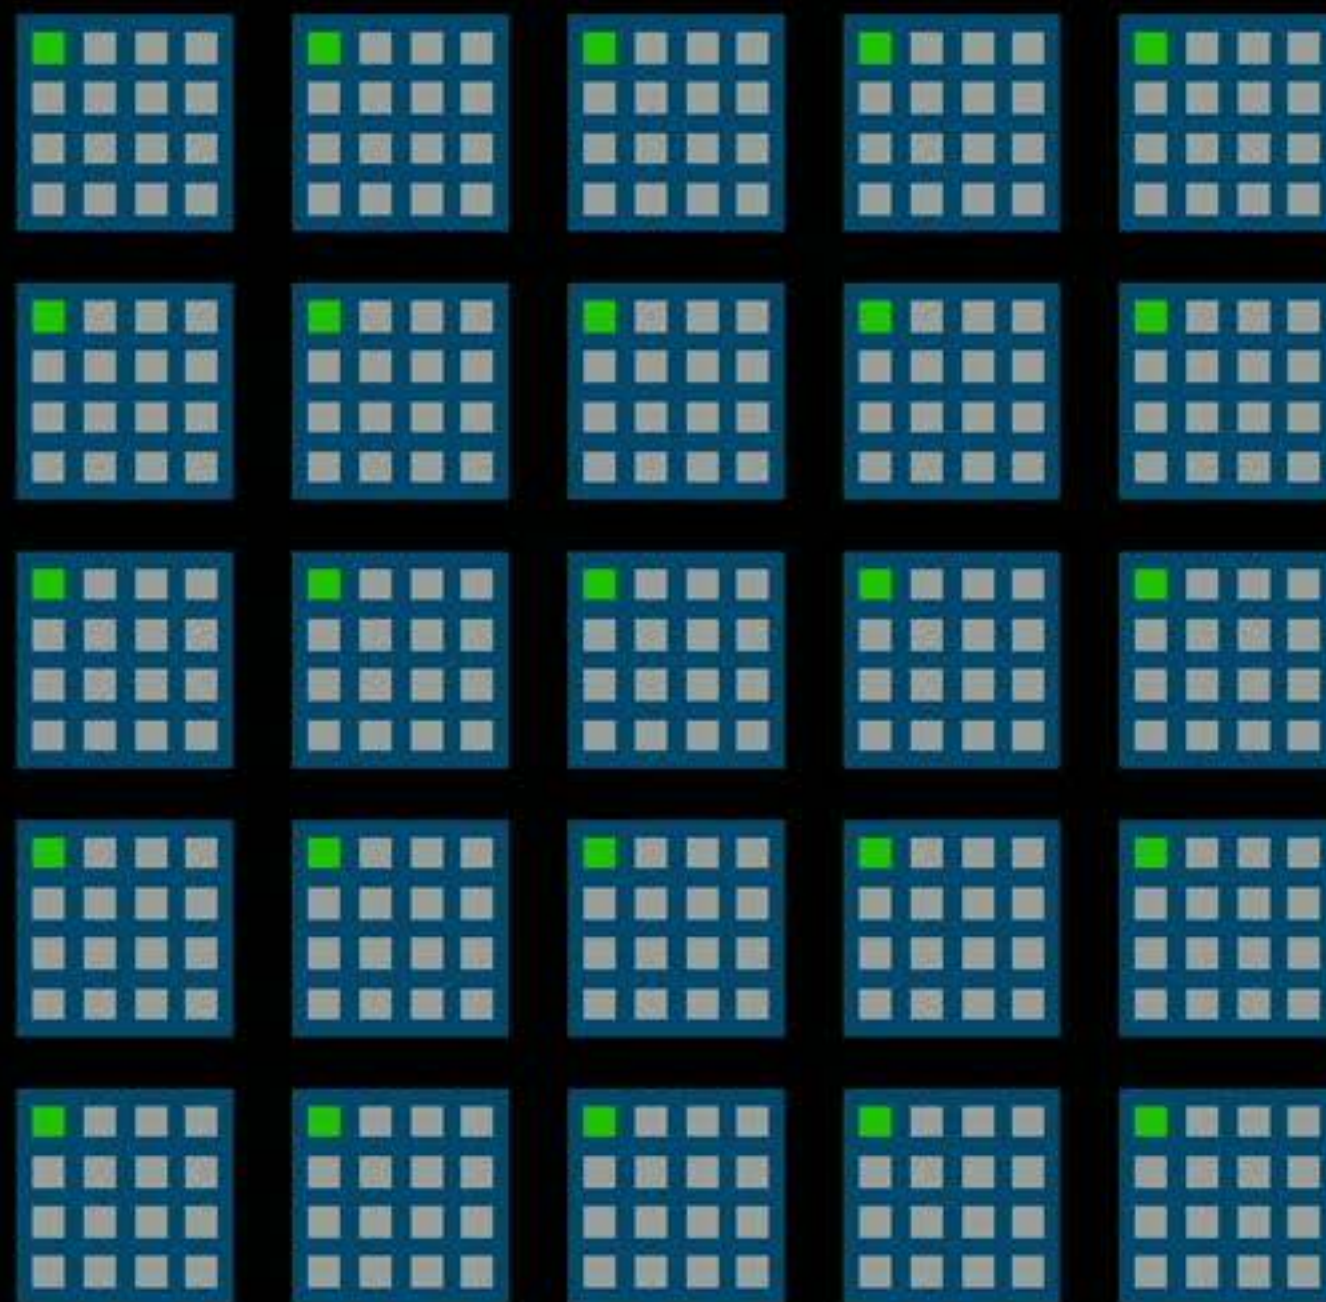
Worst Case:

many


low occupancy,

non-meshable


pages



**Standard
allocators**



`malloc(256)`



**Standard
allocators**

`malloc(256)`



**Standard
allocators**


`malloc(256)`



**Standard
allocators**

Mesh uses **randomization** to ensure live objects are uniformly distributed


Regular allocation patterns are real



How to **randomize** allocation?

Random probing:


```
while true:  
    if rand_off().is_free:  
        return rand_off
```





```
malloc(256)
```

Random probing:

```
while true:  
    if rand_off().is_free:  
        return rand_off
```



```
malloc(256)
```



Random probing:

```
while true:  
    if rand_off().is_free:  
        return rand_off
```




```
malloc(256)
```

Random probing:

```
while true:  
    if rand_off().is_free:  
        return rand_off
```

(DieHard [Berger & Zorn 2006])



Random probing fast in
expectation *iff page
occupancy is low*

**but this is at odds with
minimizing heap size!**

**Shuffle Vector:
Fast randomized
allocation + full page utilization**

Shuffle Vector: Fast randomized allocation

load

Page




Thread-local shuffle vector


Shuffle Vector: Fast randomized allocation

load

Page



Thread-local shuffle vector



All
Pages
Meshable

Finding pages to Mesh

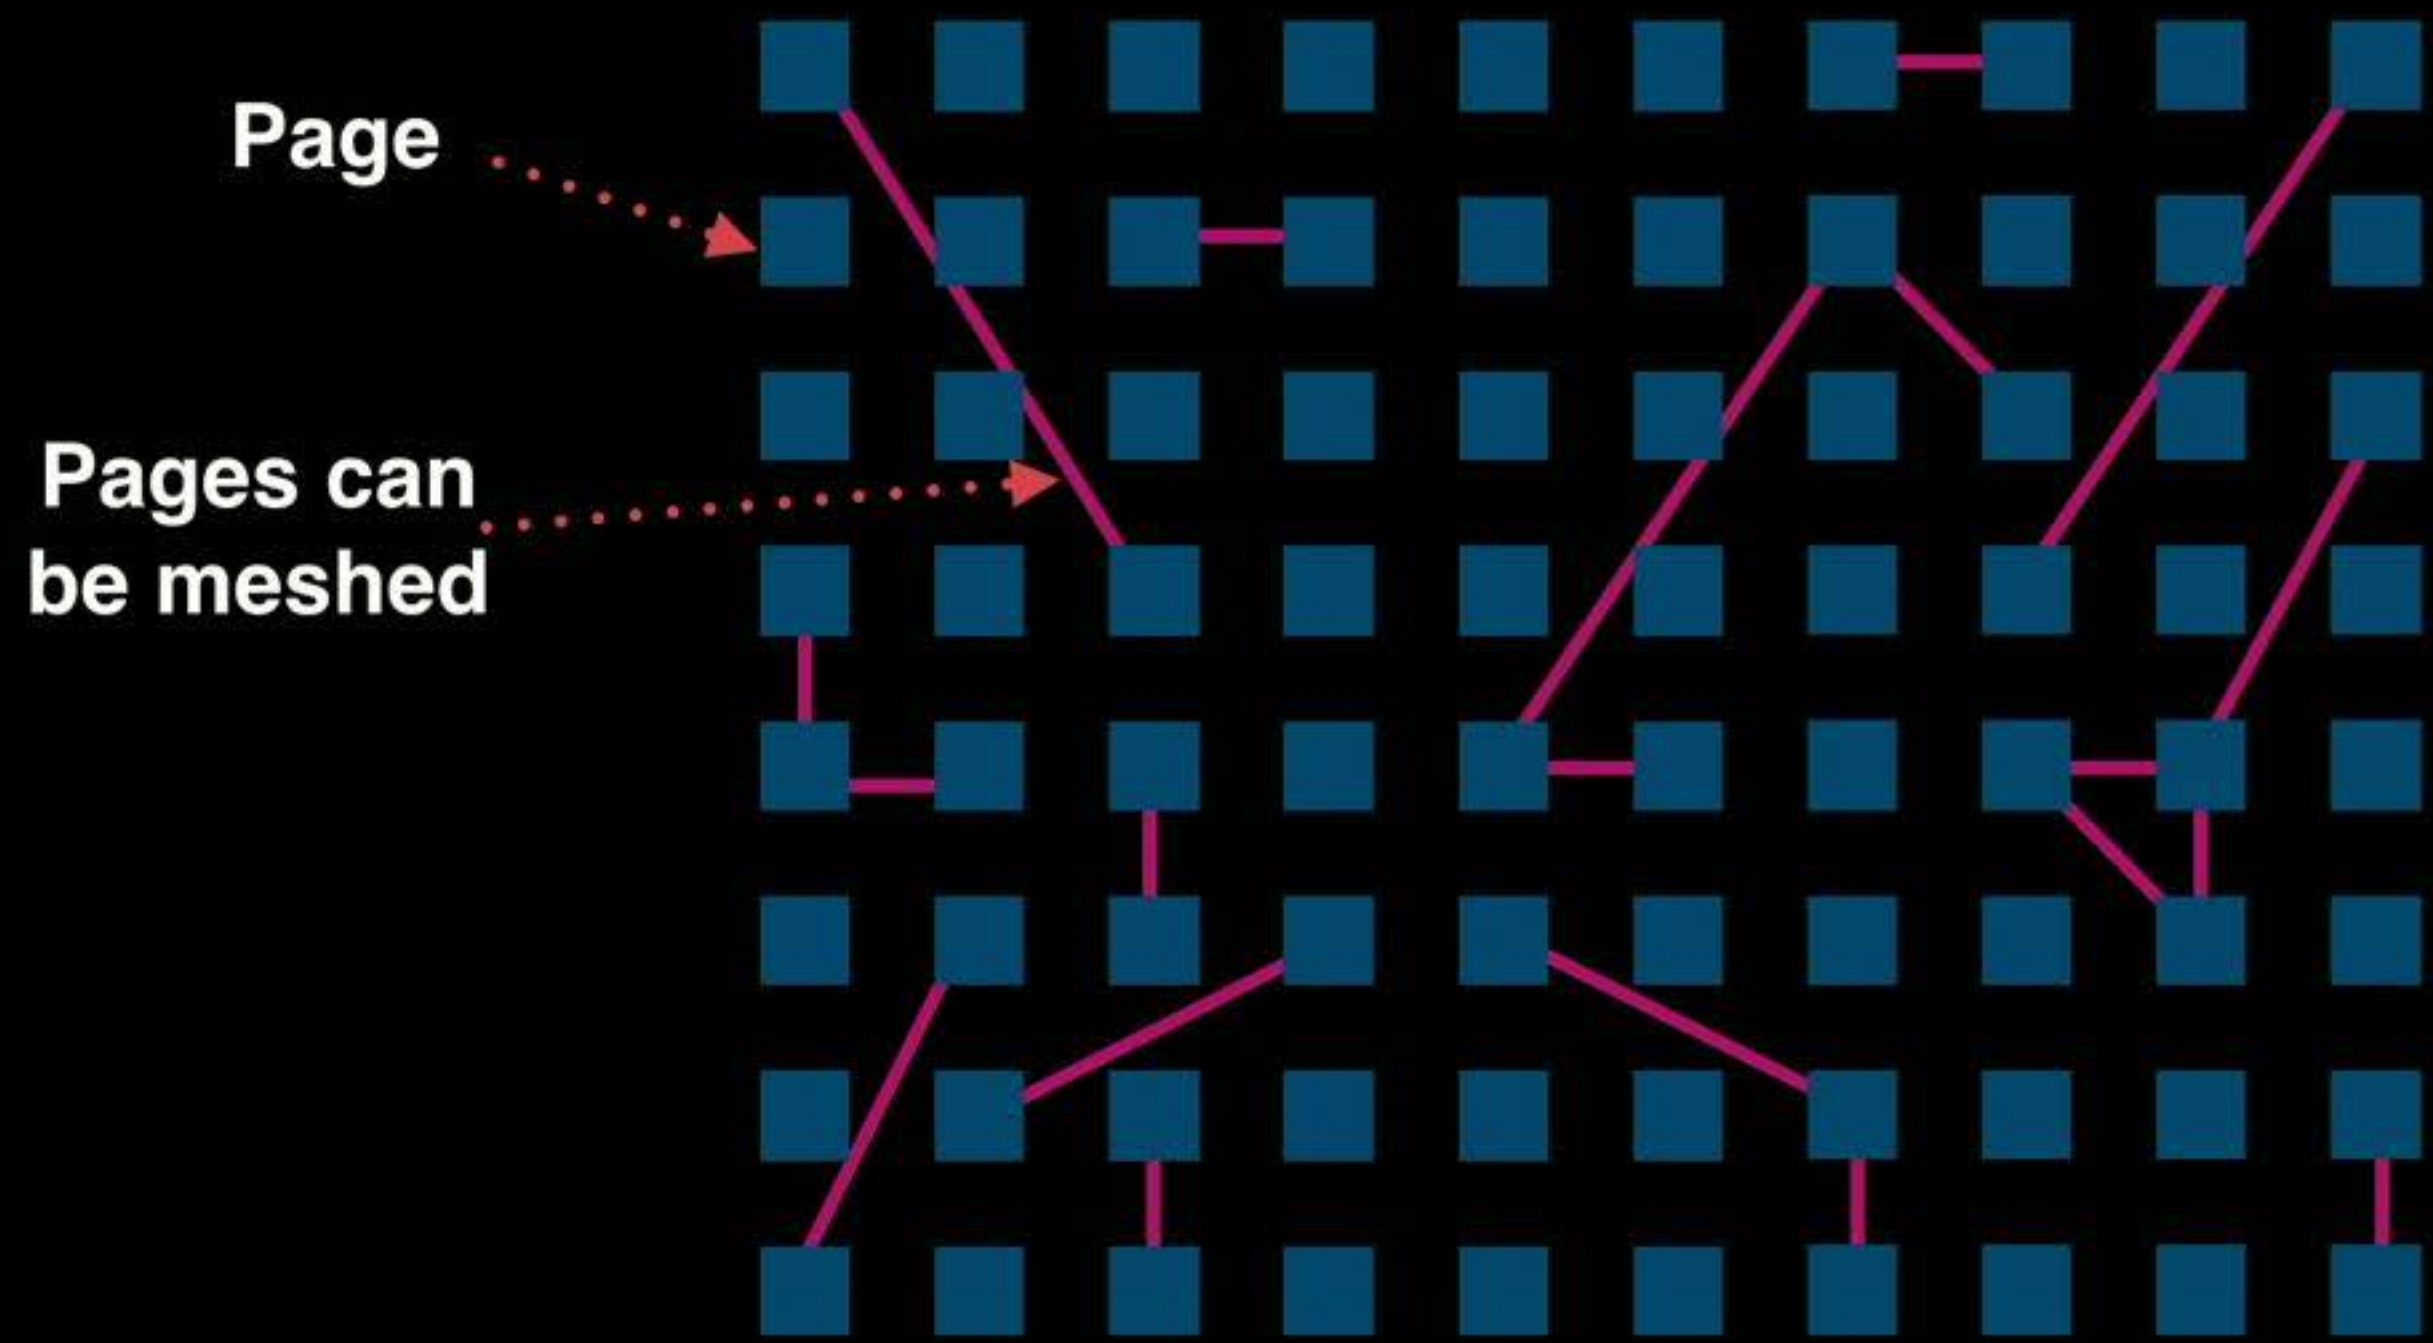
Problem: Find meshing that releases maximum number of pages

Run in the `free()` slowpath


At most once every 100 ms

Treat each size class independently

Problem: Find meshing that releases the maximum number of pages



MinCliqueCover




MinCliqueCover

(NP-Complete)


BUT! Randomness ensures we can get
away with solving simpler graph problem
(Matching)

Wrinkle: building this graph would require RAM + time




SplitMesher: approximates Matching without materializing meshing graph

Set of partially full pages



Pairs of meshable pages







Iterate, comparing $a[i]$ to $b[i]$




loop, comparing $a[i]$ to $b[(i+1)\%len]$






Remove found match



Continue



SplitMesher: approximates Matching
without materializing meshing graph

$O(n/q)$ time

(q is the global probability of spans meshing)

SplitMesher: approximates Matching
without materializing meshing graph

$O(n/q)$ time


(q is the global probability of spans meshing)

$1/2^*$ approximation w.h.p.



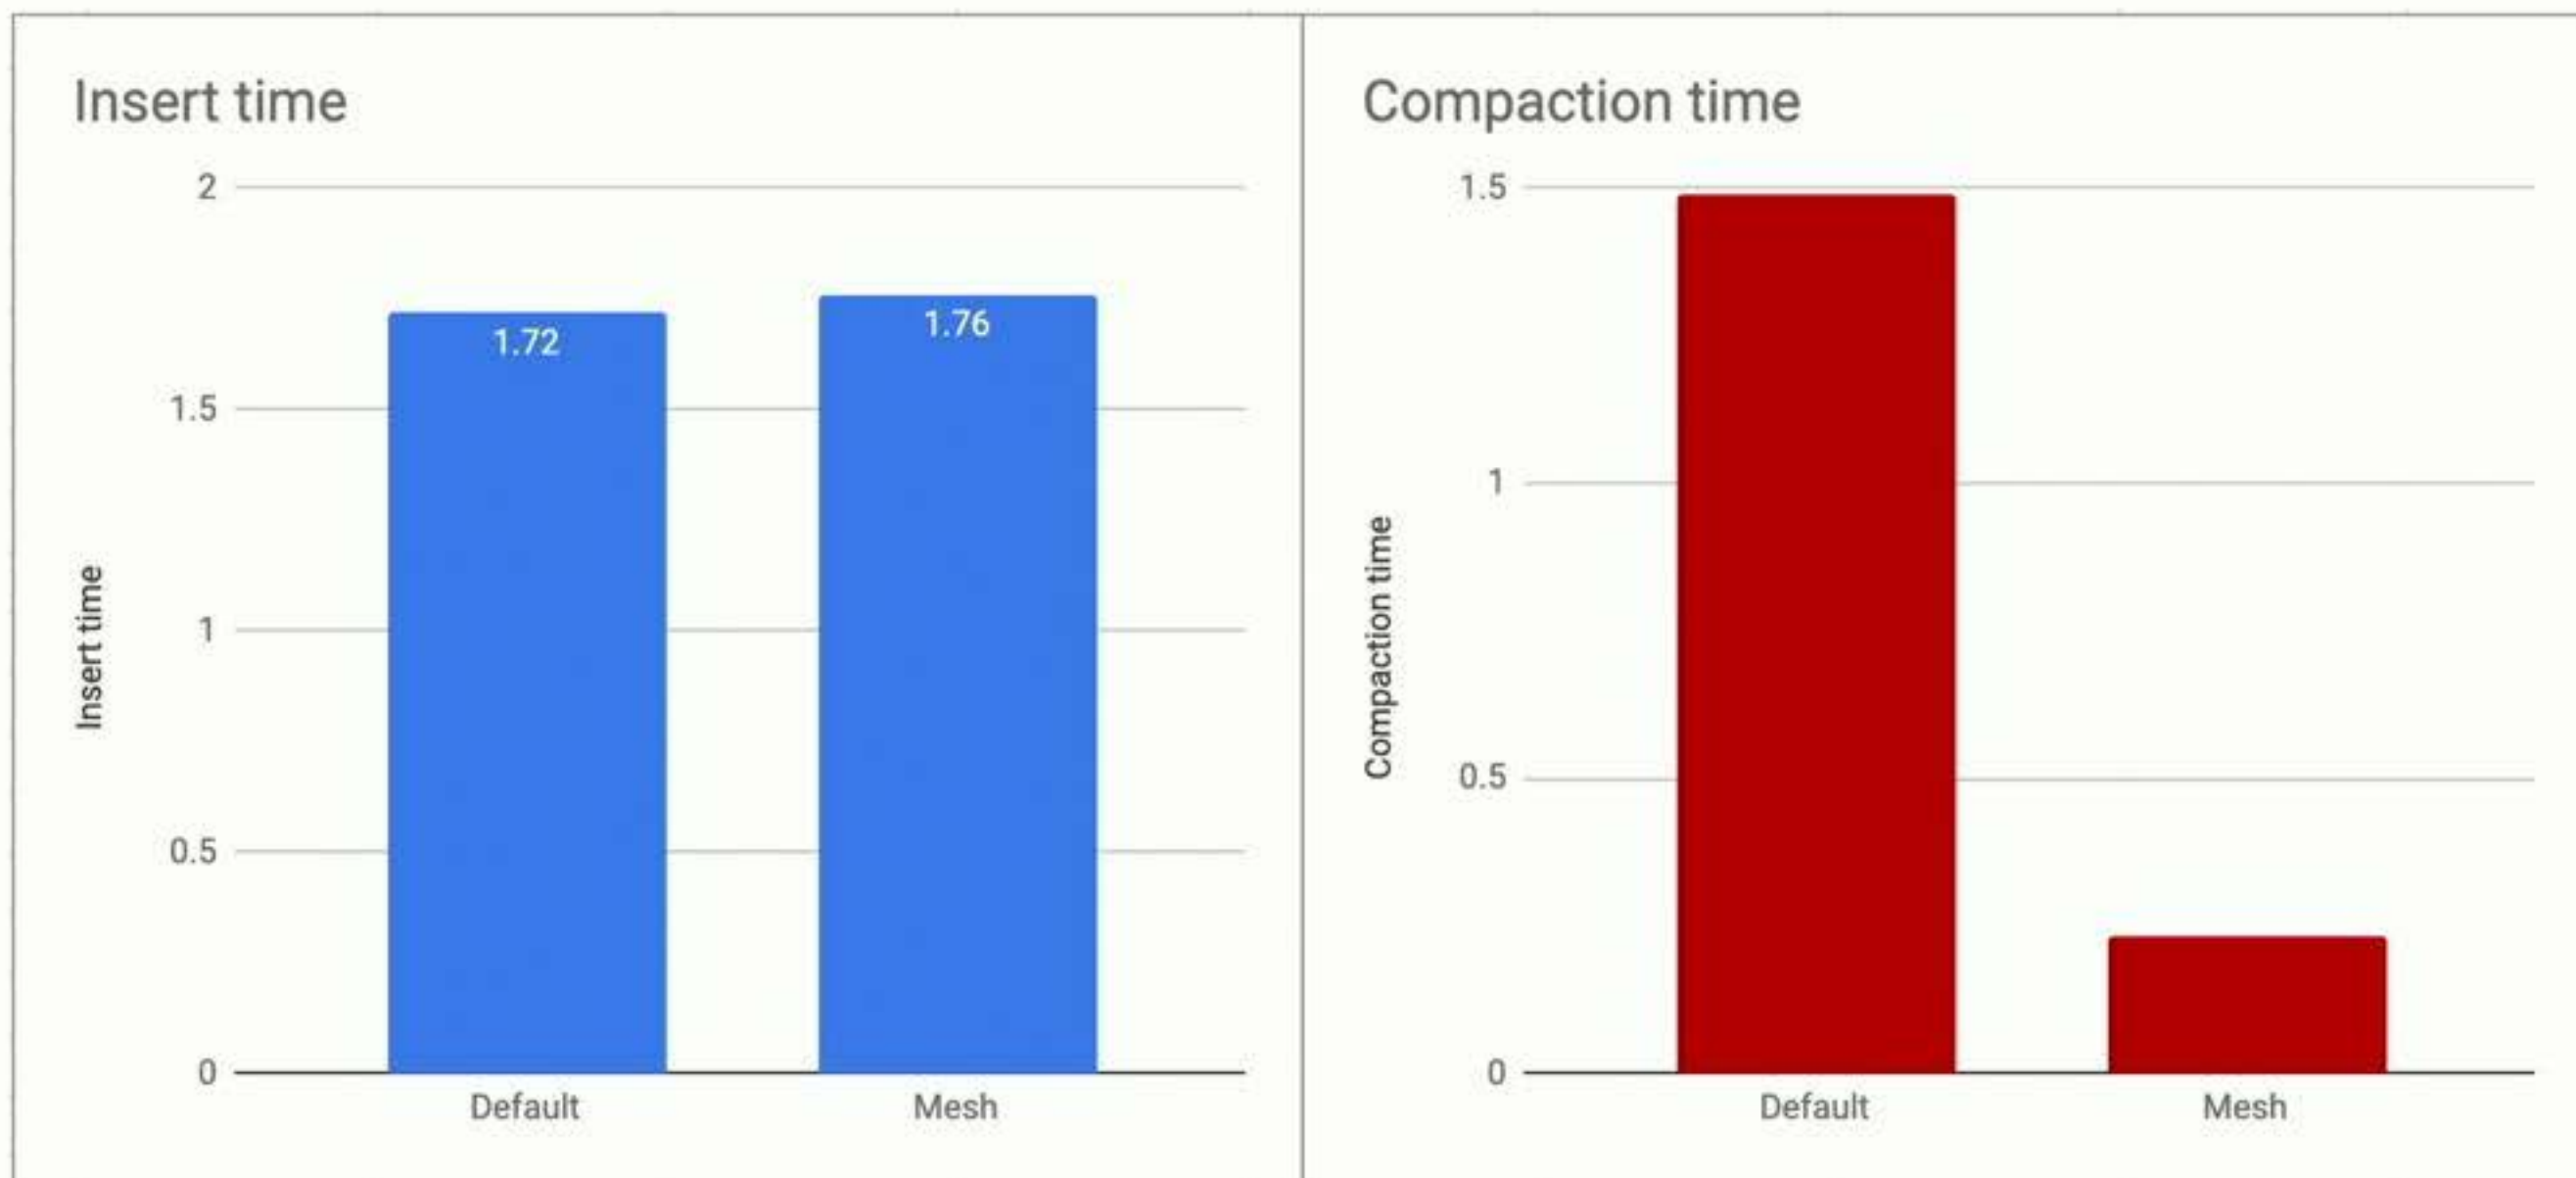
redis

- jemalloc + activedefrag
- Mesh
- no compaction






redis + MESH




17% heap size reduction

< 1% performance overhead



—default jemalloc—Mesh

Ruby Compaction for Free



[http:// LIBMESH .org](http://LIBMESH.org)

¡Compacción sin Relocación!
(compaction without relocation)

