

Research Topics in Database Management

Bigger, Faster, and Stronger Systems

Niv Dayan

Who am I?

> 12 years of research experience in data structures & algorithms for databases

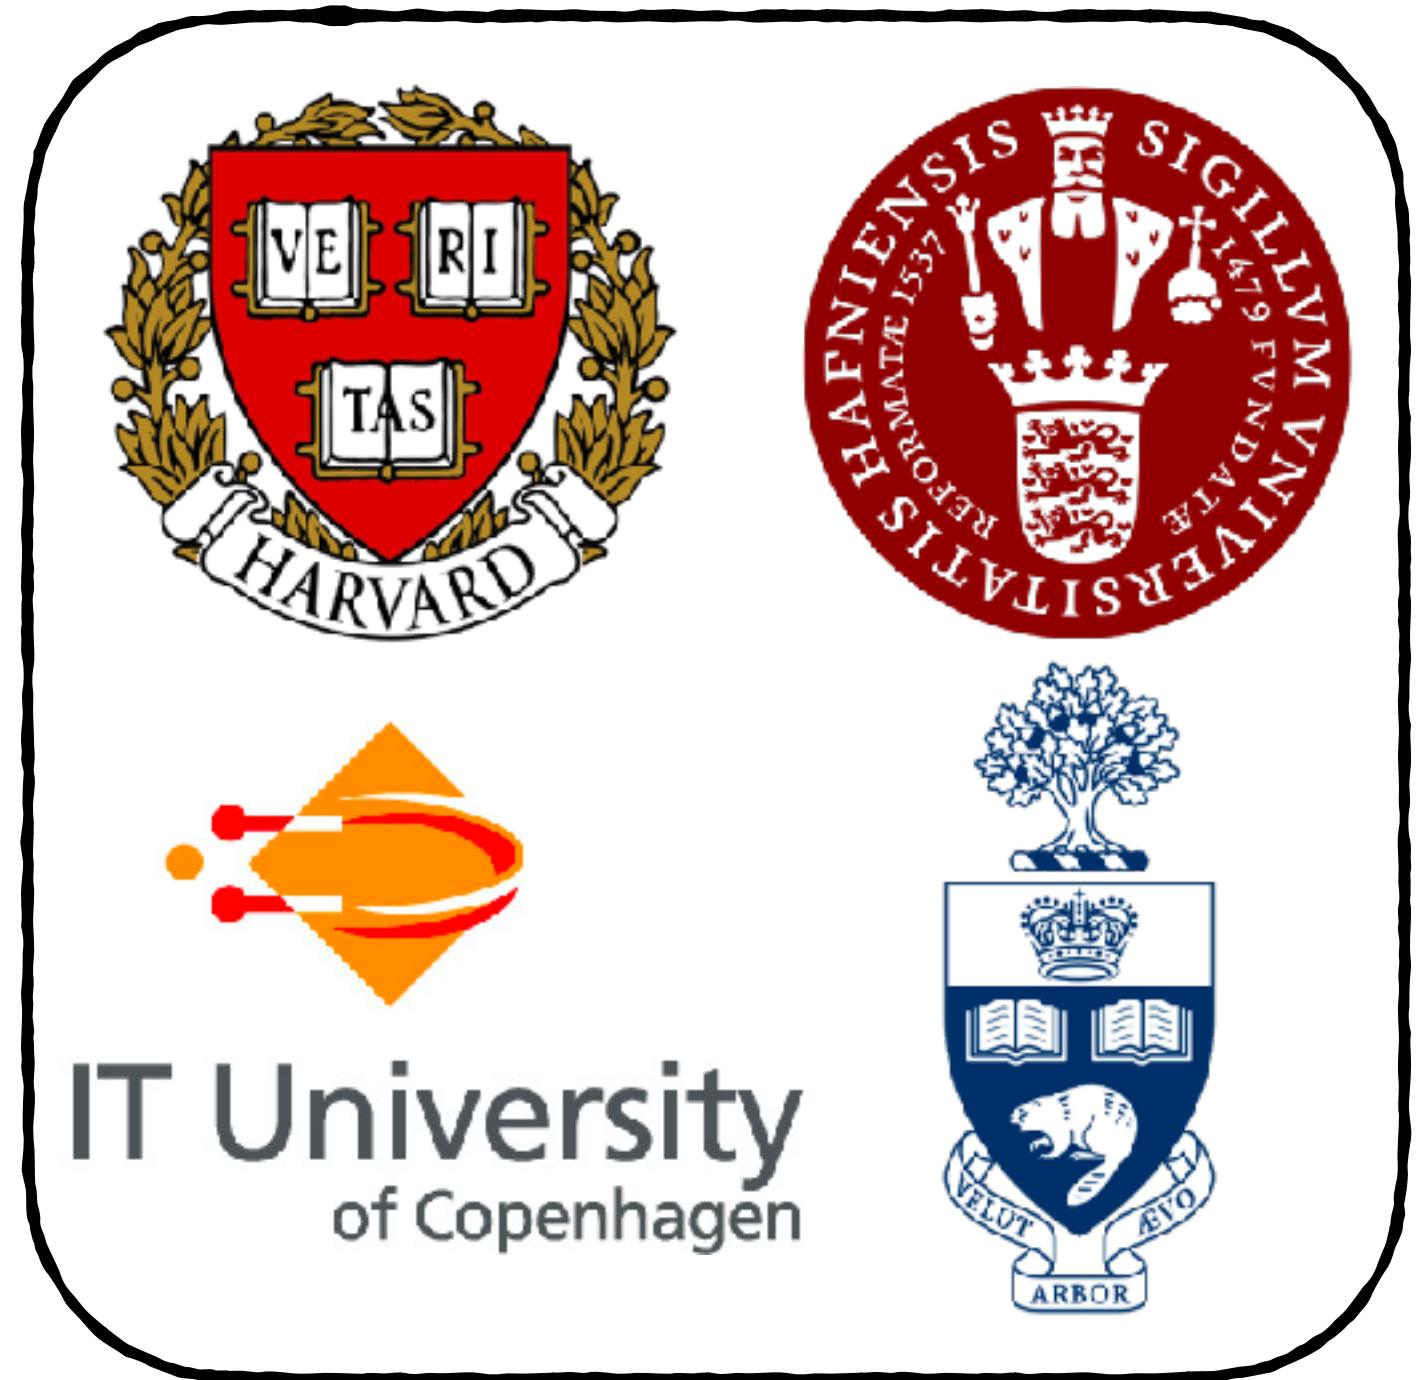


<https://www.nivdayan.net/>

Who am I?

> 12 years of research experience in data structures & algorithms for databases

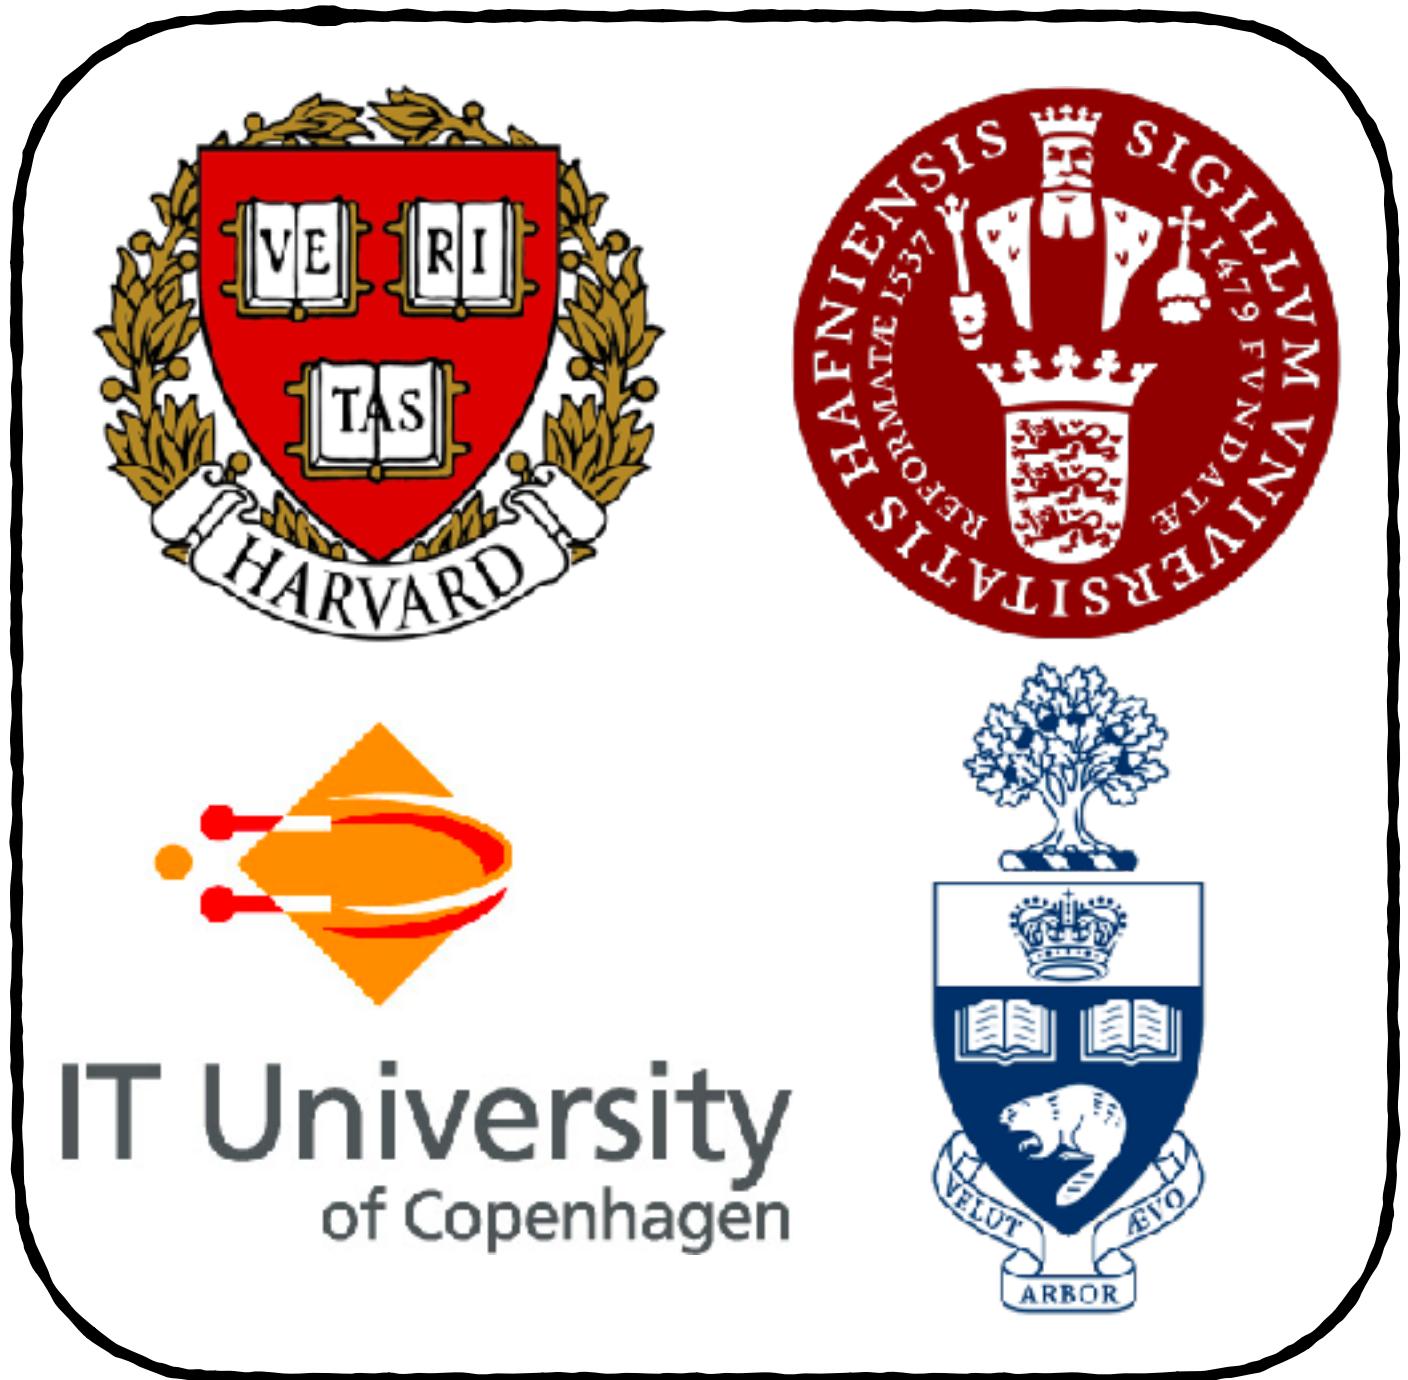
In academia



Who am I?

> 12 years of research experience in data structures & algorithms for databases

In academia



And in industry

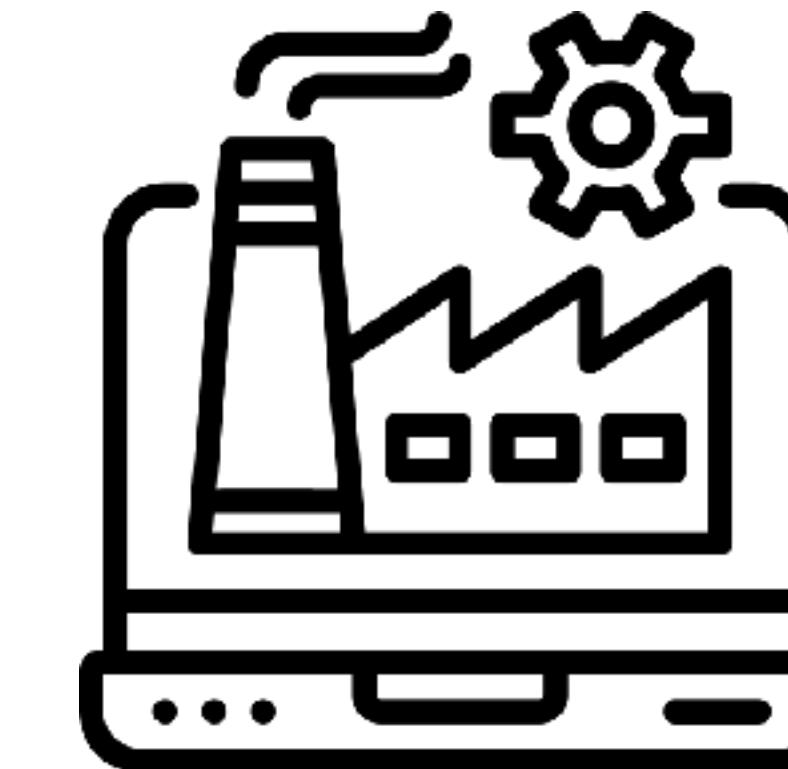


This course combines both

Theory



Practice



For data structures & algorithms for databases

Who are you?



Who are you?

Undergrad

Grad

Who are you? Prerequisites...

Who are you? Prerequisites...

Operating Systems

Concurrency & synchronization
File systems, virtual memory

Who are you? Prerequisites...

Operating Systems

Concurrency & synchronization
File systems, virtual memory

Design and Analysis of Data Structures

Binary trees, sorting, hash tables, priority queues, Big-O analysis

Who are you? Prerequisites...

Operating Systems

Concurrency & synchronization
File systems, virtual memory

Design and Analysis of Data Structures

Binary trees, sorting, hash tables, priority
queues, Big-O analysis

Database Internals e.g., (CSC443)

Storage, buffer pools, B-trees, transactions,
write-ahead logging, query processing, etc.

Who are you? Prerequisites...

Operating Systems

Concurrency & synchronization
File systems, virtual memory

Design and Analysis of Data Structures

Binary trees, sorting, hash tables, priority
queues, Big-O analysis

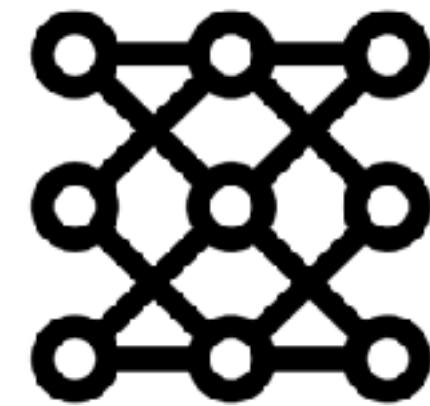
Database Internals e.g., (CSC443)

Storage, buffer pools, B-trees, transactions,
write-ahead logging, query processing, etc.

Solid programming skills in C, C++, Java, or at least Python

Background Knowledge

**CSC443 is
background for some
topics**



**All lectures
recorded**



**Will let you
know what to
catch up on**

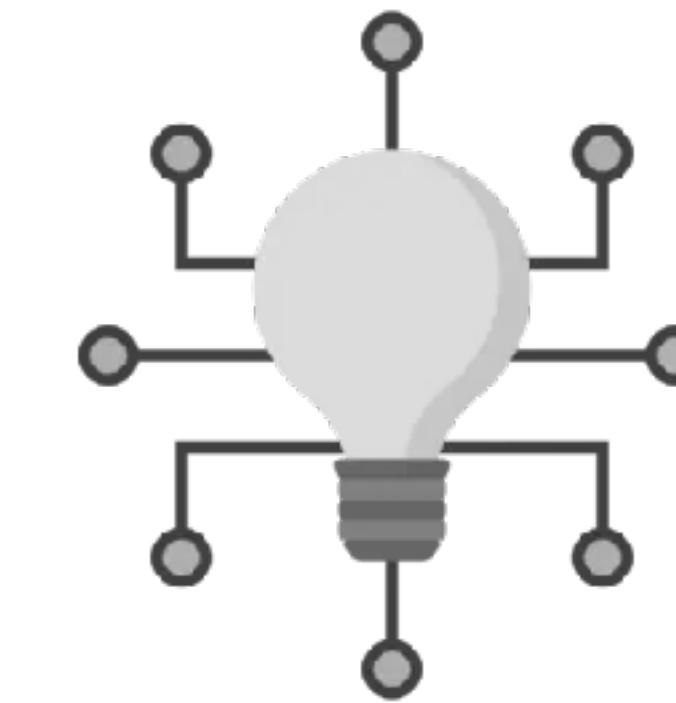


<https://www.nivdayan.net/database-system-technology-csc443>

Data Structures Seminar



**Reading ≈20-30
Papers**

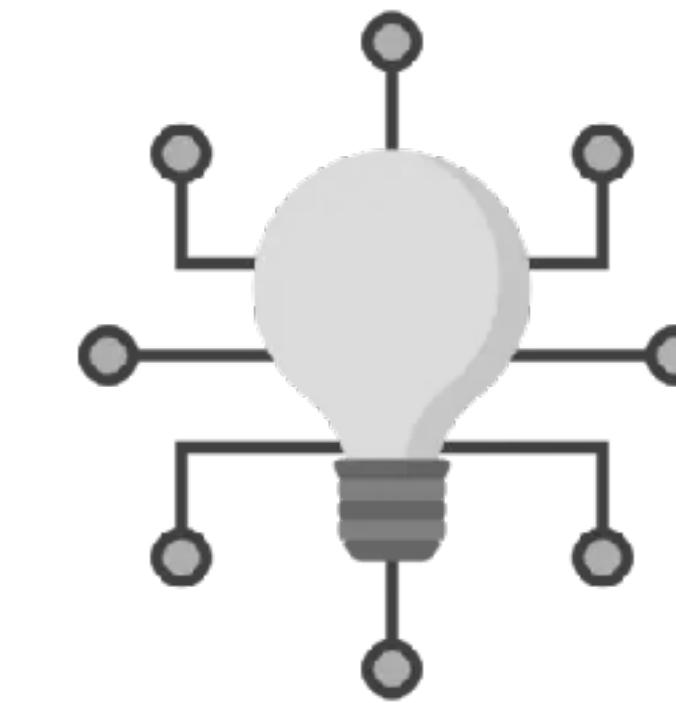


**Small Research
Project**

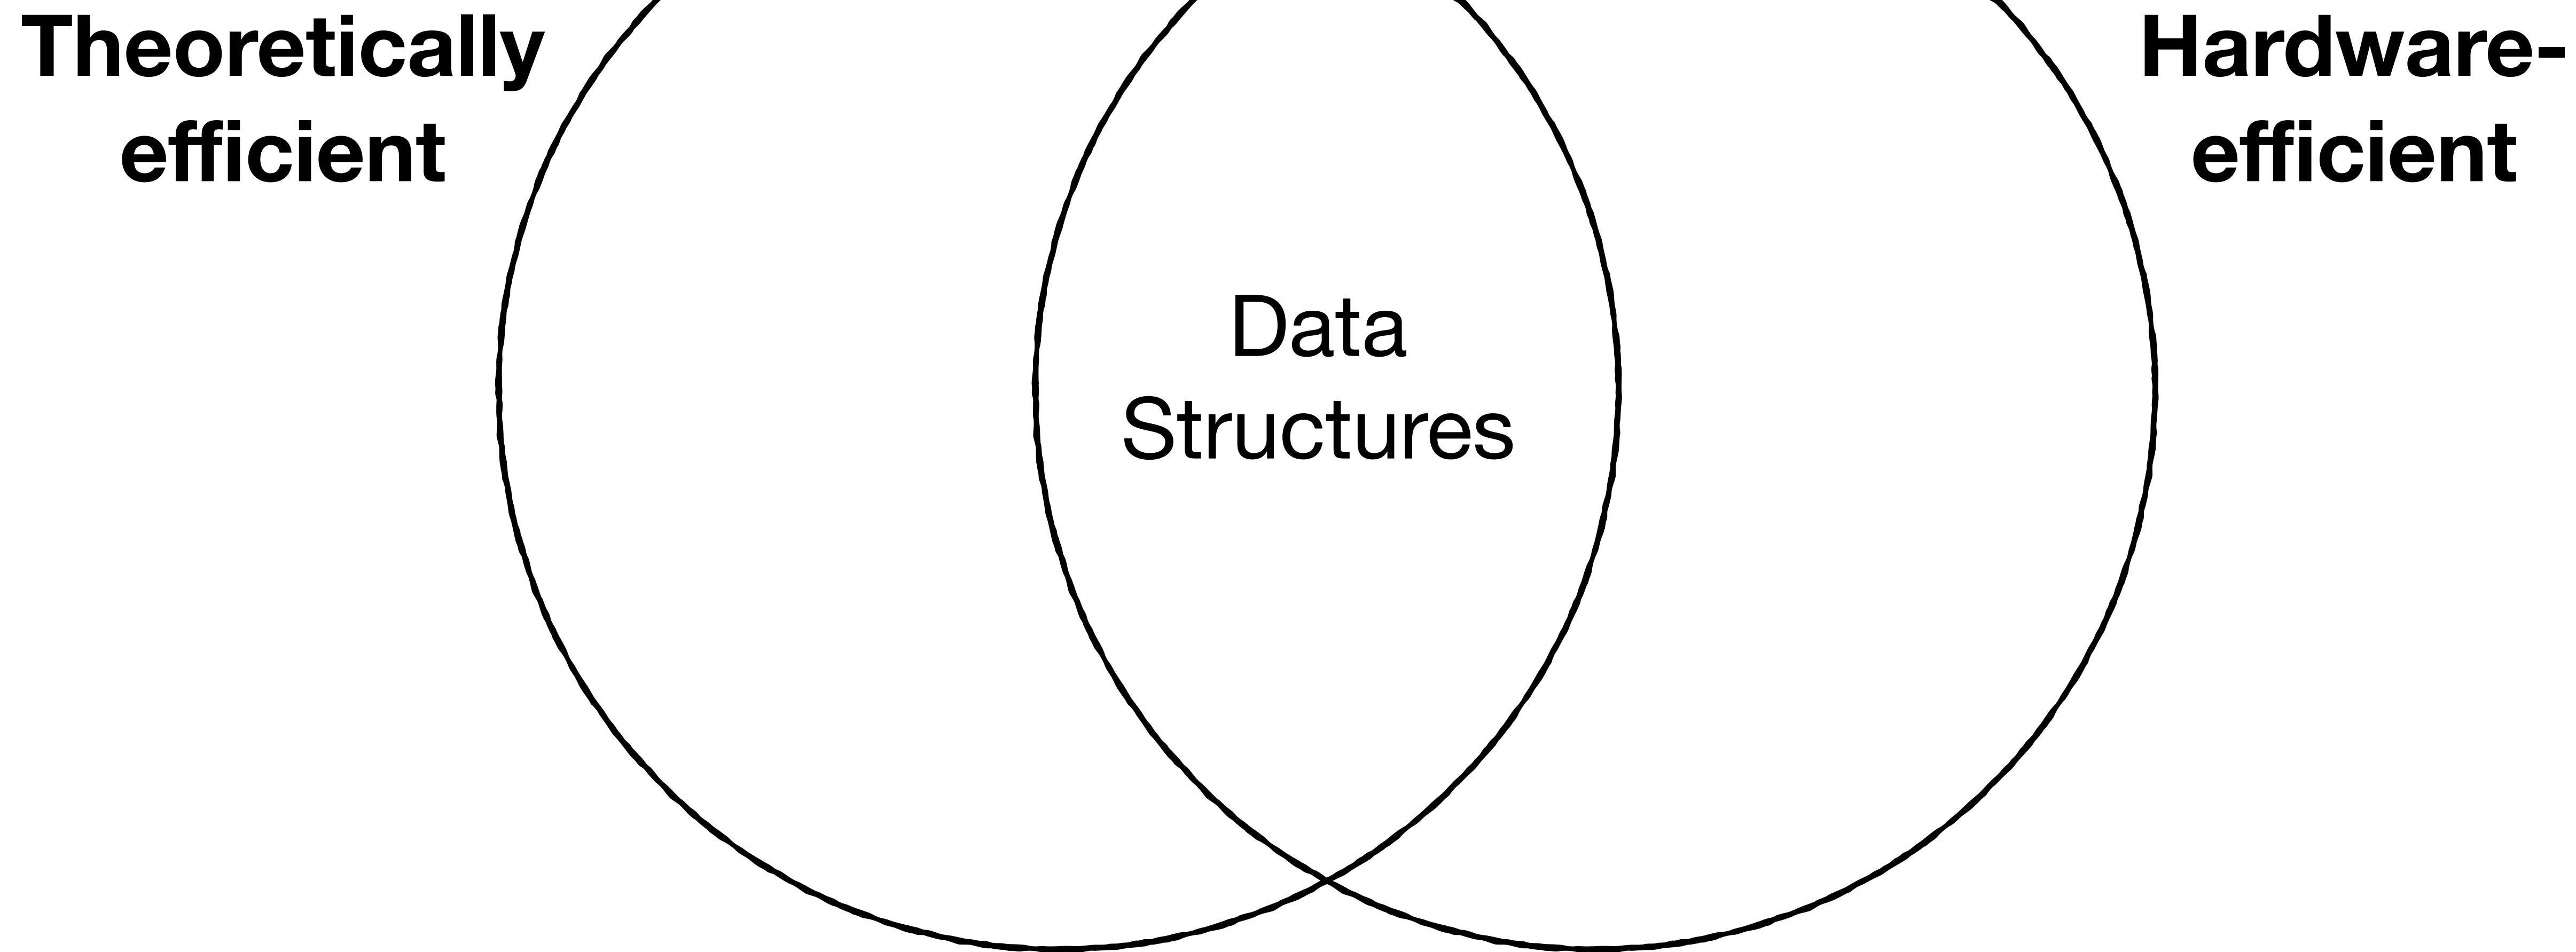
Data Structures Seminar



Reading ≈20-30
Papers



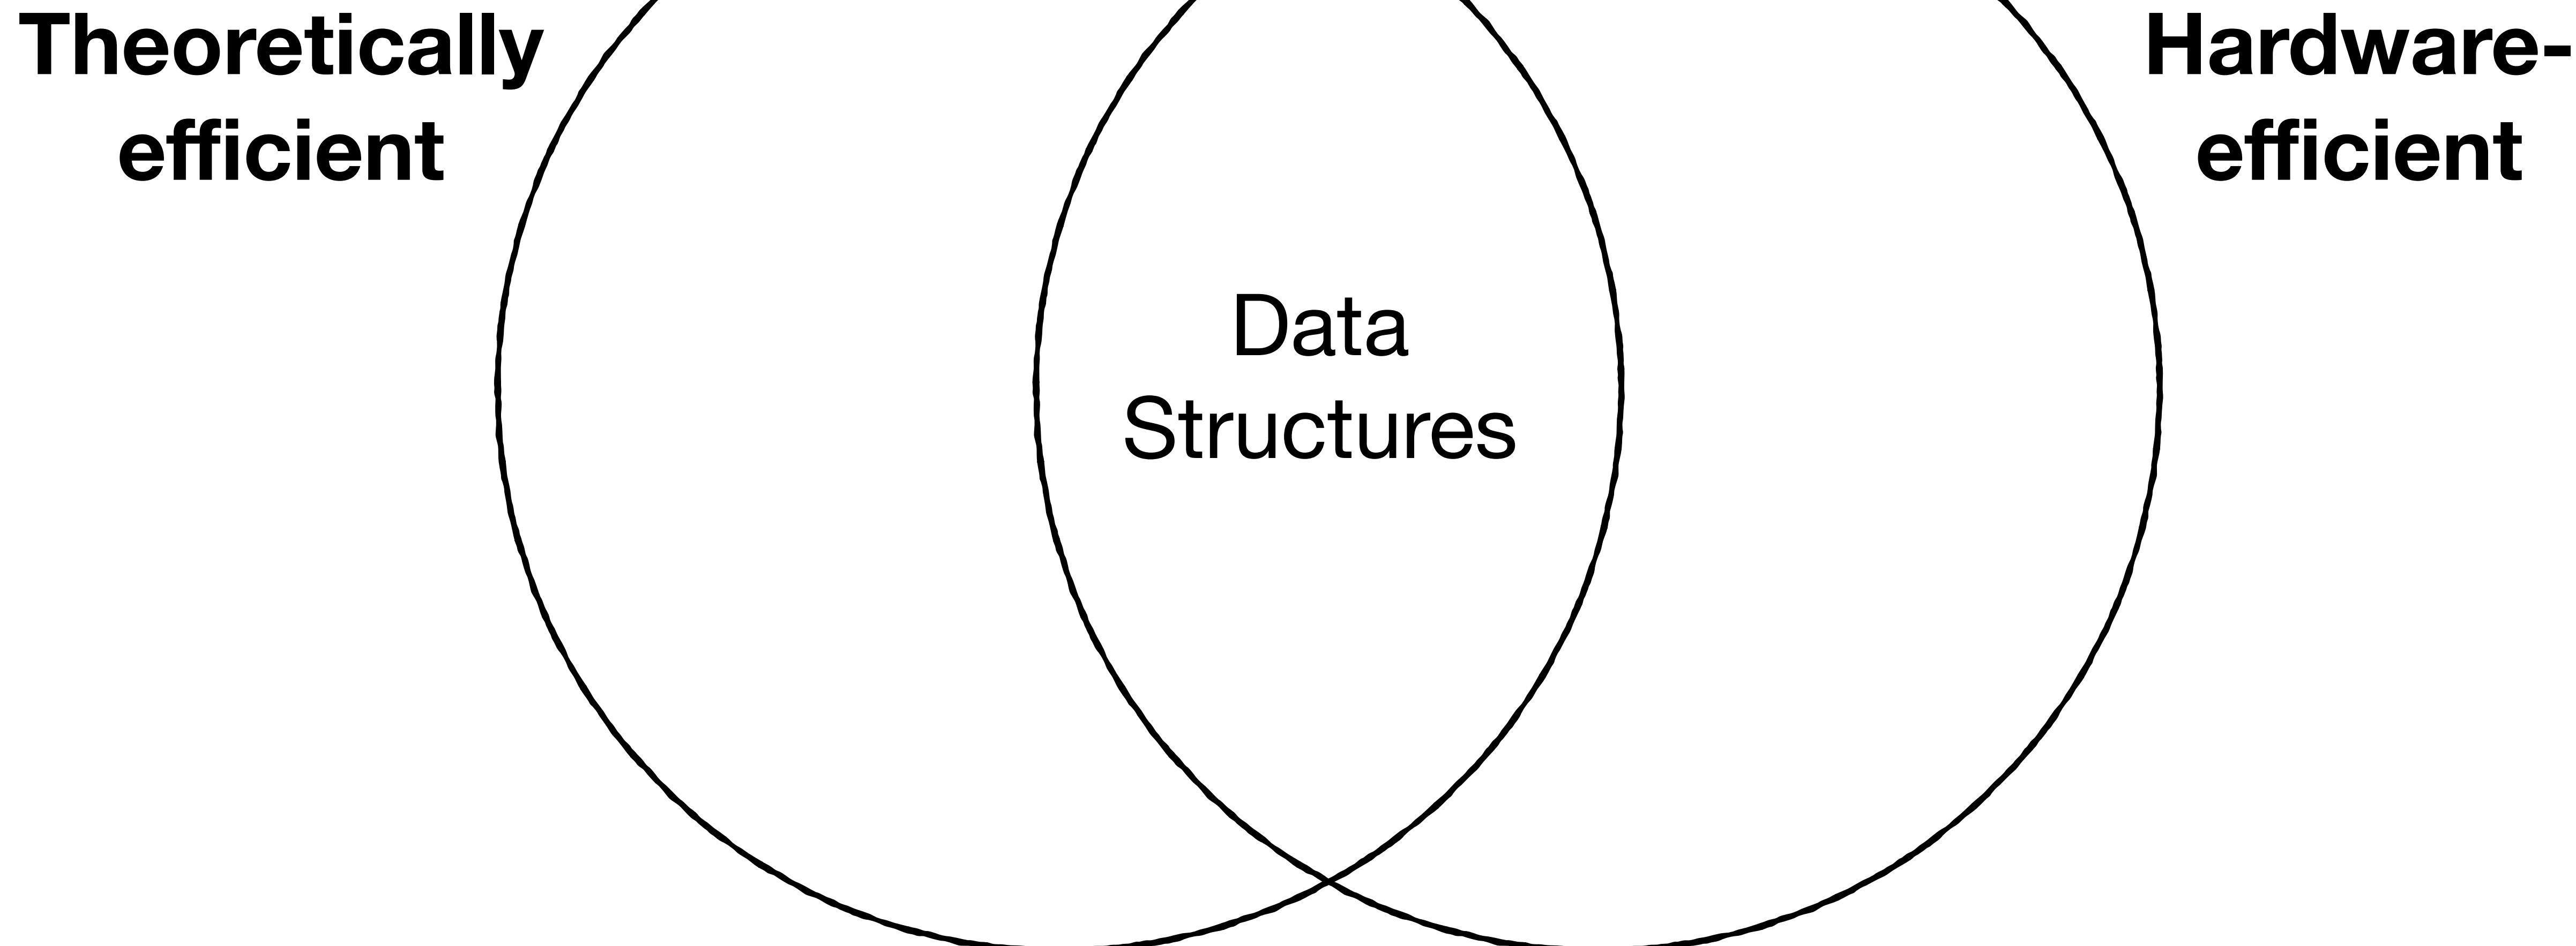
**Small Research
Project**



**Theoretically
efficient**

**Hardware-
efficient**

Data
Structures



Theoretically
efficient

Hardware-
efficient

Data
Structures

Important for your maturity as engineers/researchers who
can achieve high performance

Why read papers?

Reading papers
is a skill



Get research
ideas



Employ the state of
the art



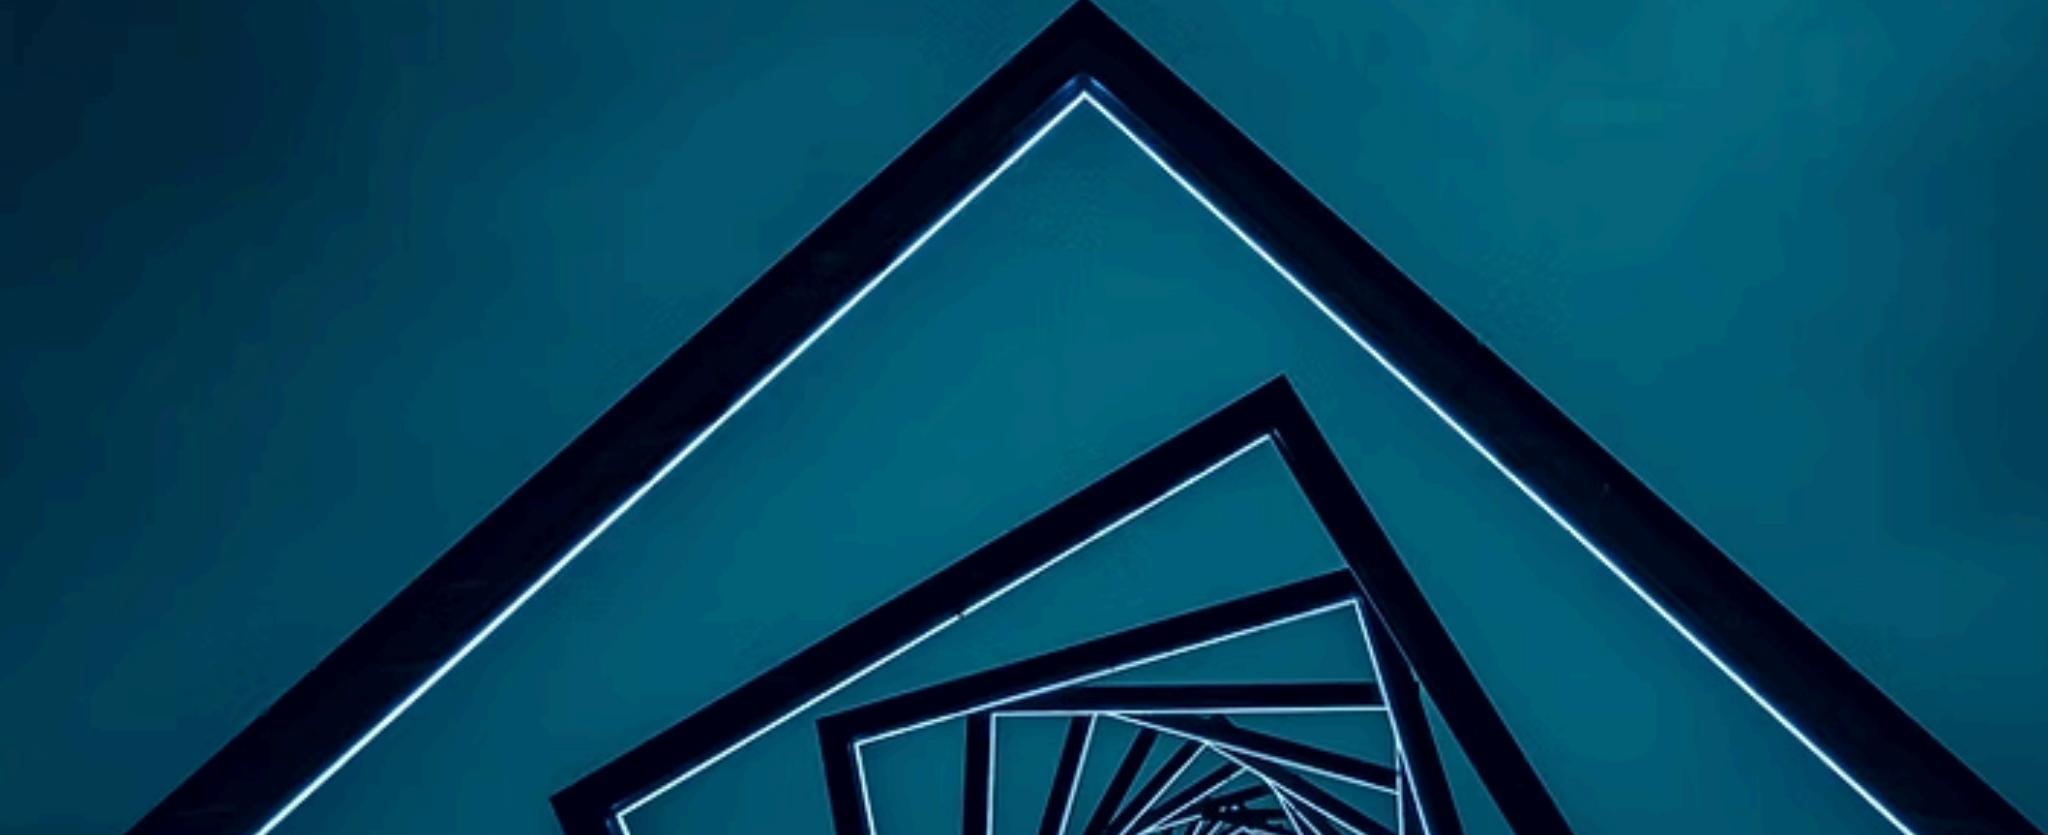
Website



CSC2525

Research Topics in Database Management

Instructor: Niv Dayan
Lectures: Wednesday 15:00-17:00 (UC 85)
Office Hours: after each class



<https://www.nivdayan.net/research-topics-in-database-management-csc2525>

12 Class Sessions



Use the first two lectures wisely

**Dynamic Arrays &
Filter Data
Structures**

Enjoy the material?
you're in the right
place



Use the first two lectures wisely

Dynamic Arrays &
Filter Data
Structures

**Enjoy the material?
you're in the right
place**



Participation

**You are required to
attend each class**



**Read papers in
advance**

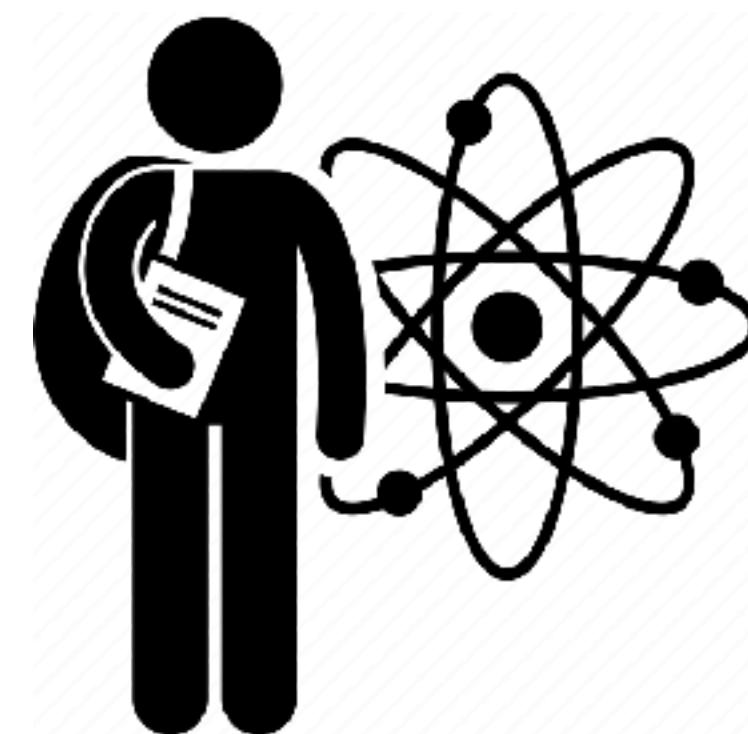


**Participate in
class discussions**



Project

**Implement &
evaluate**

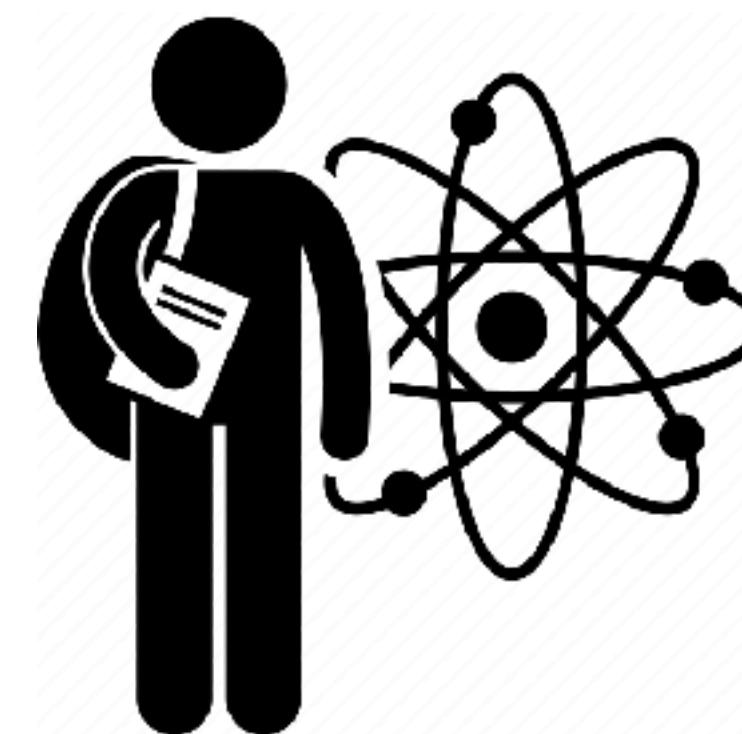


**Proposals due by
mid-Feb**



Project

**Implement &
evaluate**



**Proposals due by
mid-Feb**

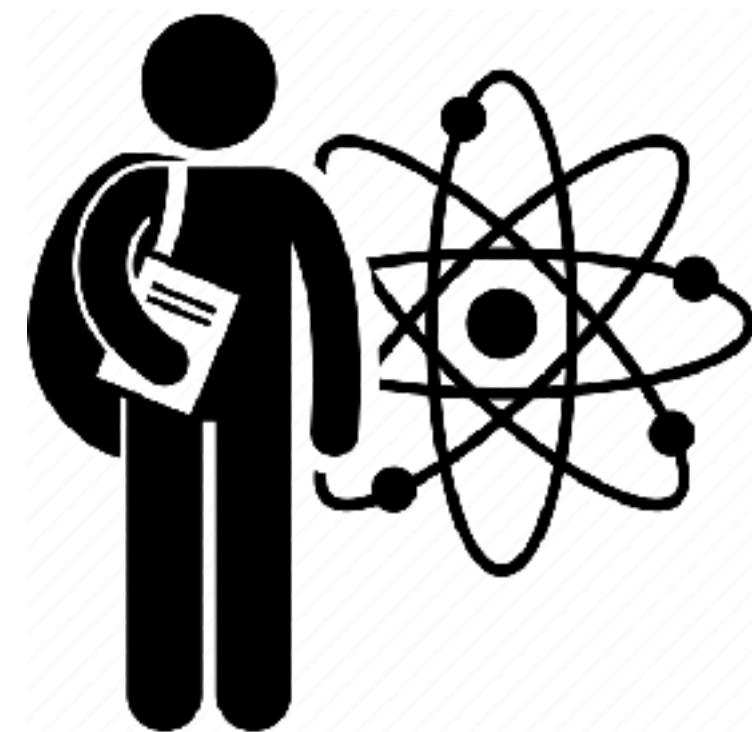


**You may start
earlier**



Project

Implement &
evaluate



Proposals due by
mid-Feb



You may start
earlier



More on this later

Written Exam

Likely 2 hours



**Likely April 7-8 or 30
(Before/after exam period)**



Grade Components

(1) Project Report & Code

(2) Oral exam

Grade Components

(1) Project Report & Code

(2) Oral exam

Precise breakdown to be announced later

Office Hours

Right after
class





Post questions for everyone's benefit!



We'll record classes, but you must still attend.

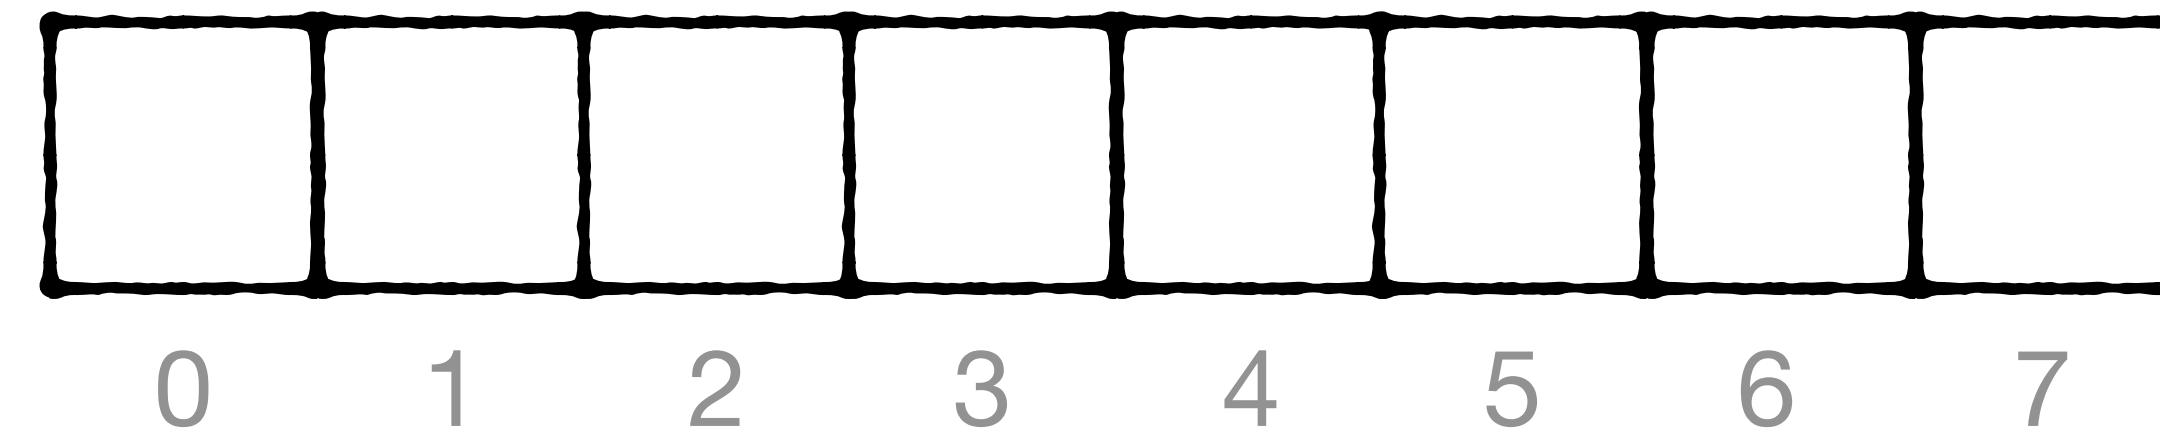
And now to our first lecture

Dynamic Arrays

CSC2525 Research Topics in Database Management

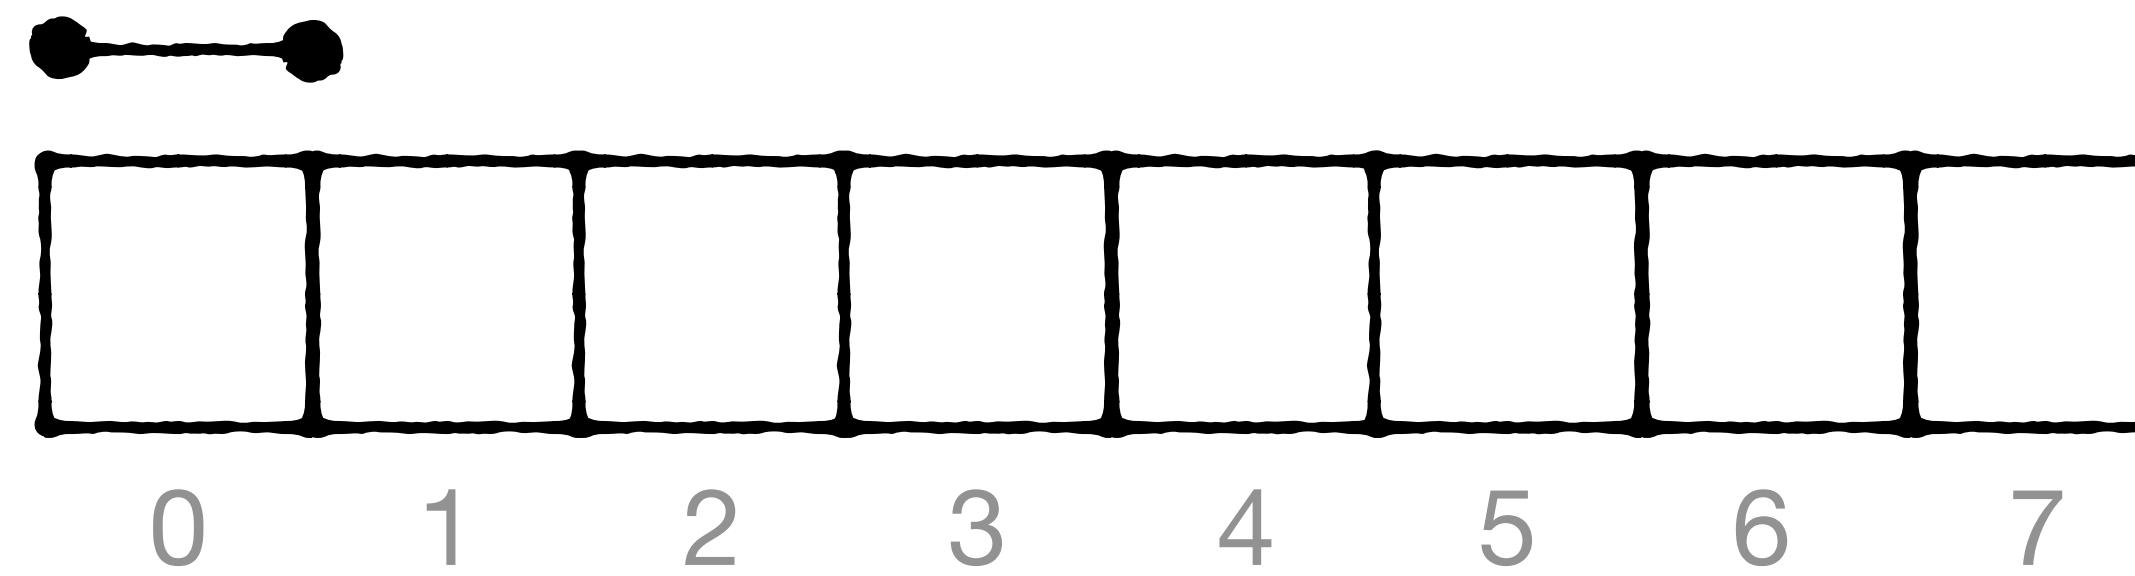
Niv Dayan

Arrays



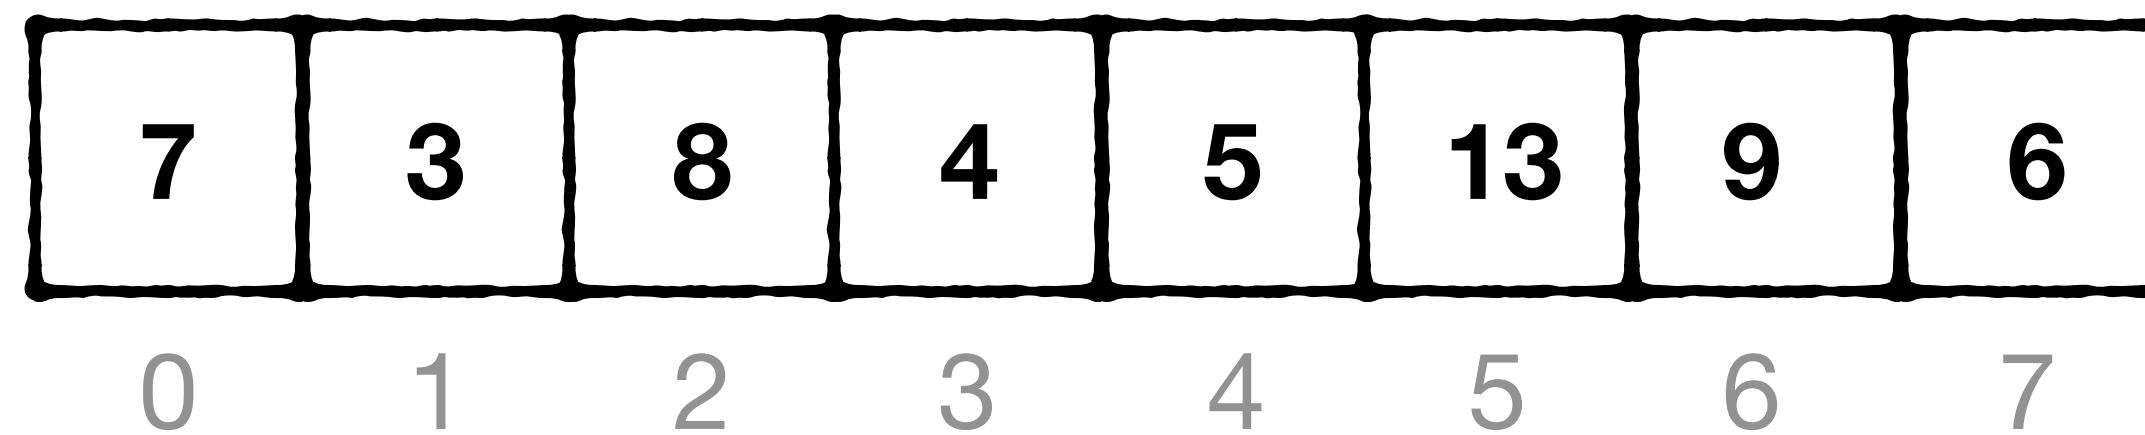
Arrays

Fixed width slots

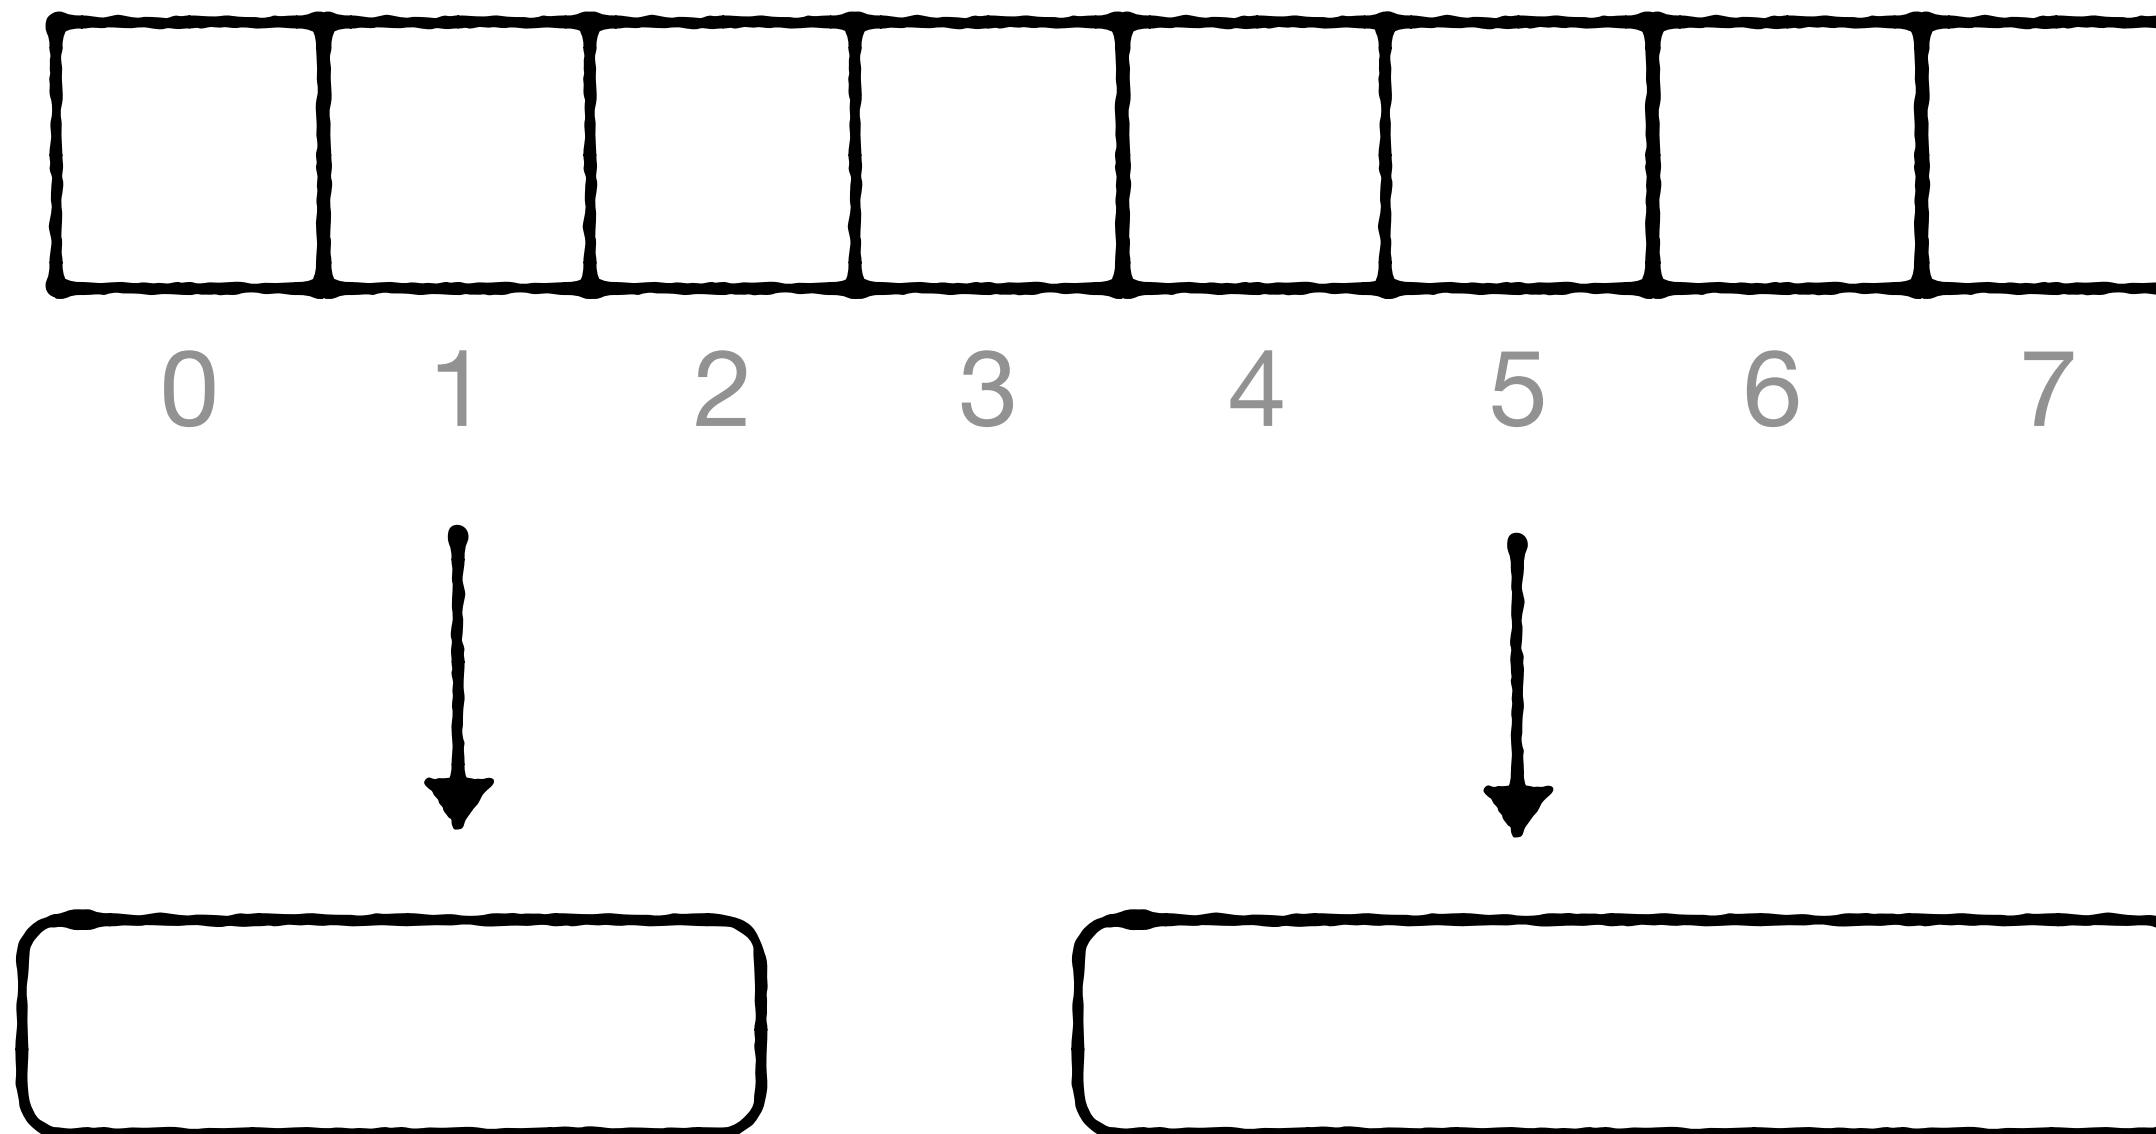


Fixed width slots

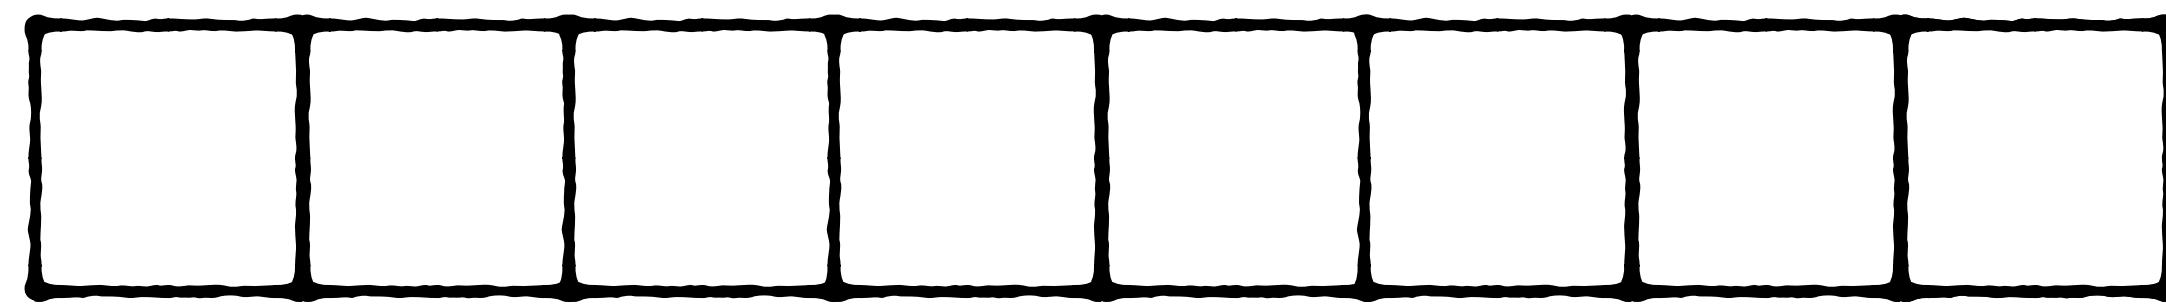
e.g., integers or floating points



Fixed width slots
Or pointers to complex types

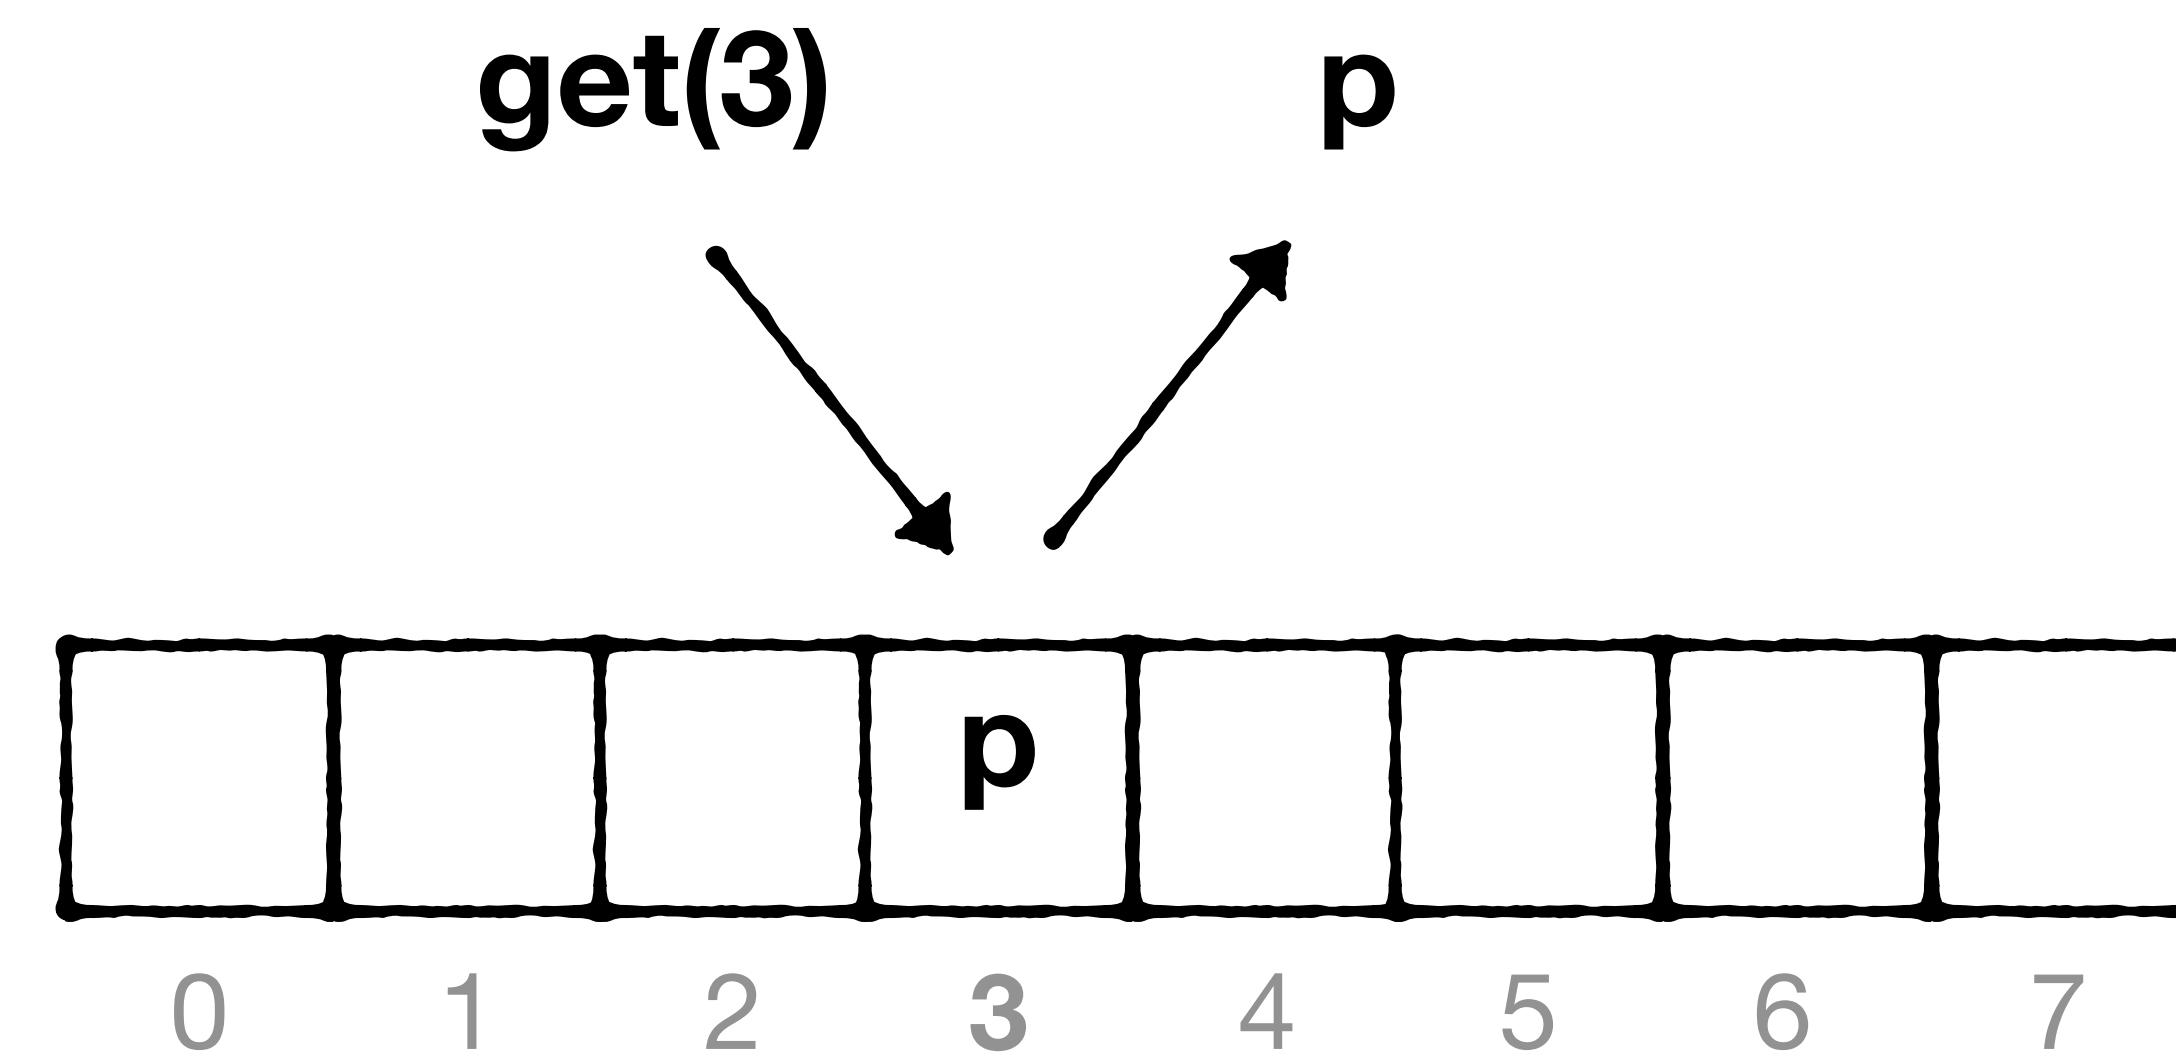


put(3, p)



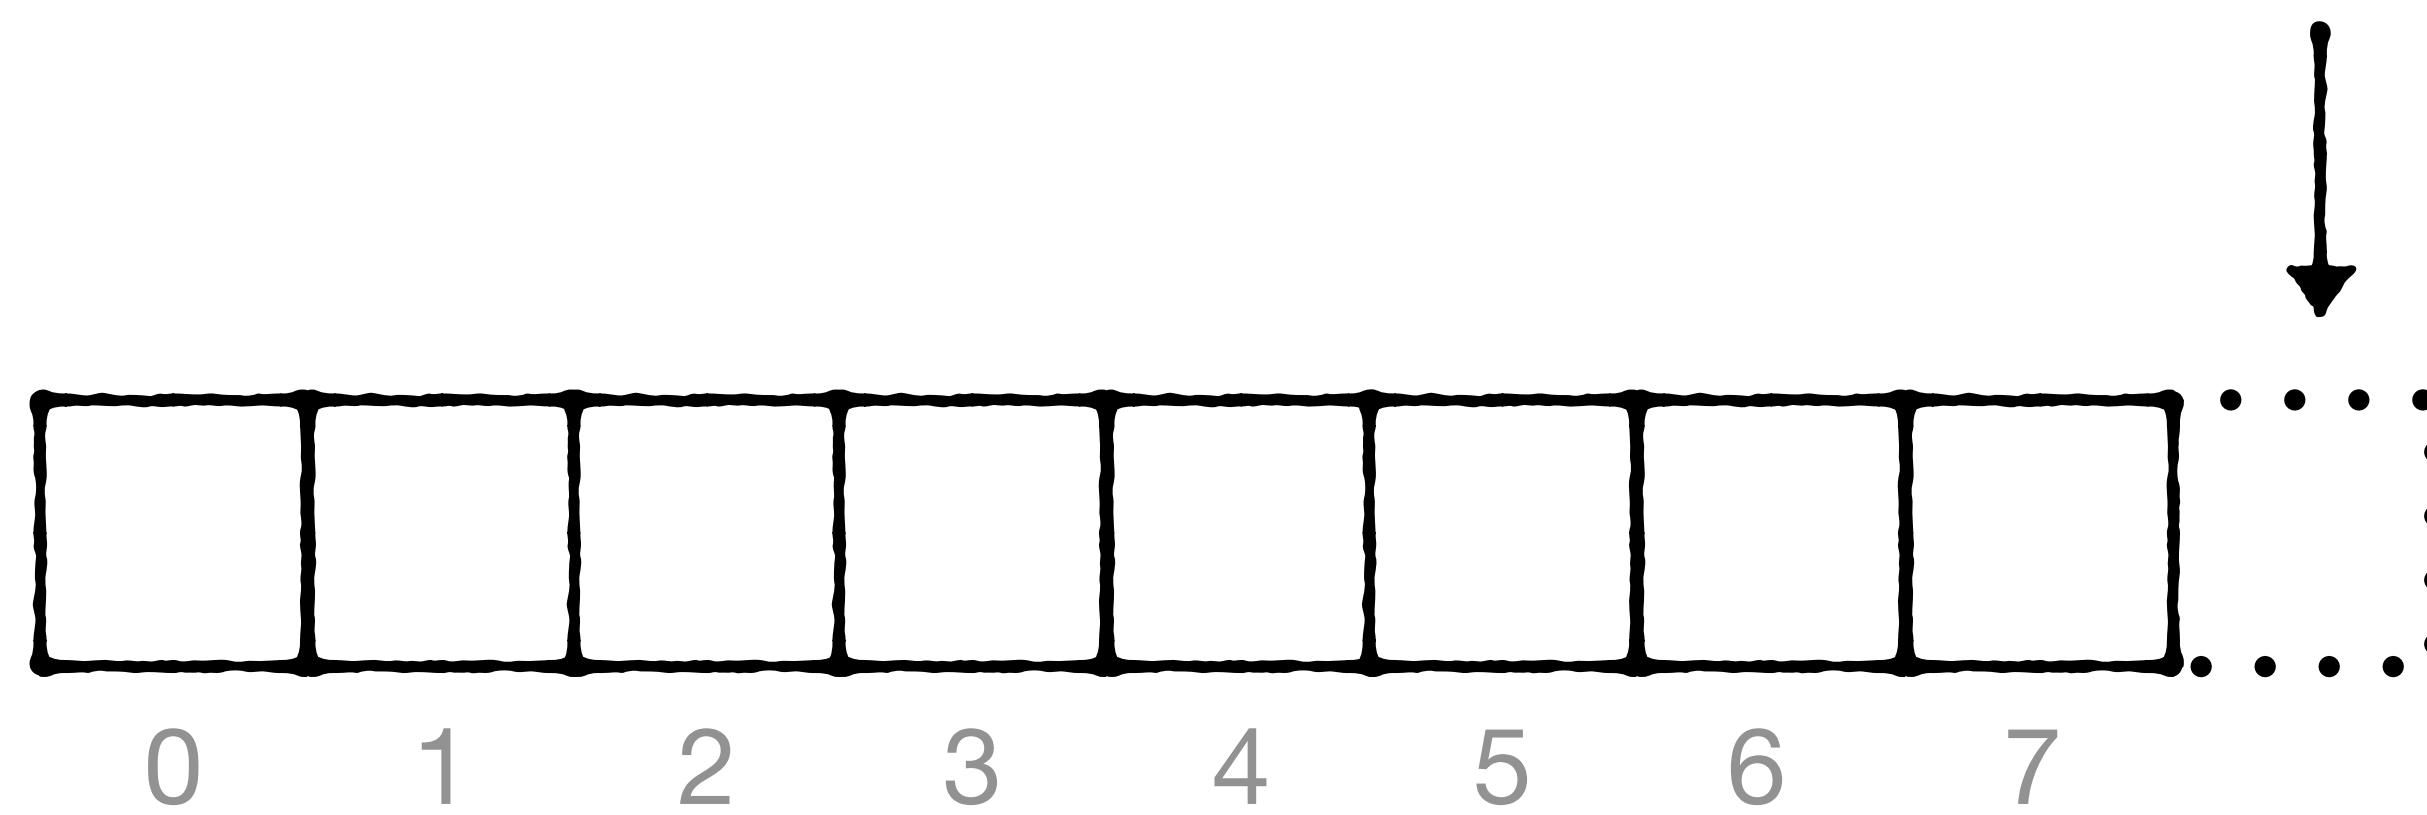
0 1 2 3 4 5 6 7

Supports random access

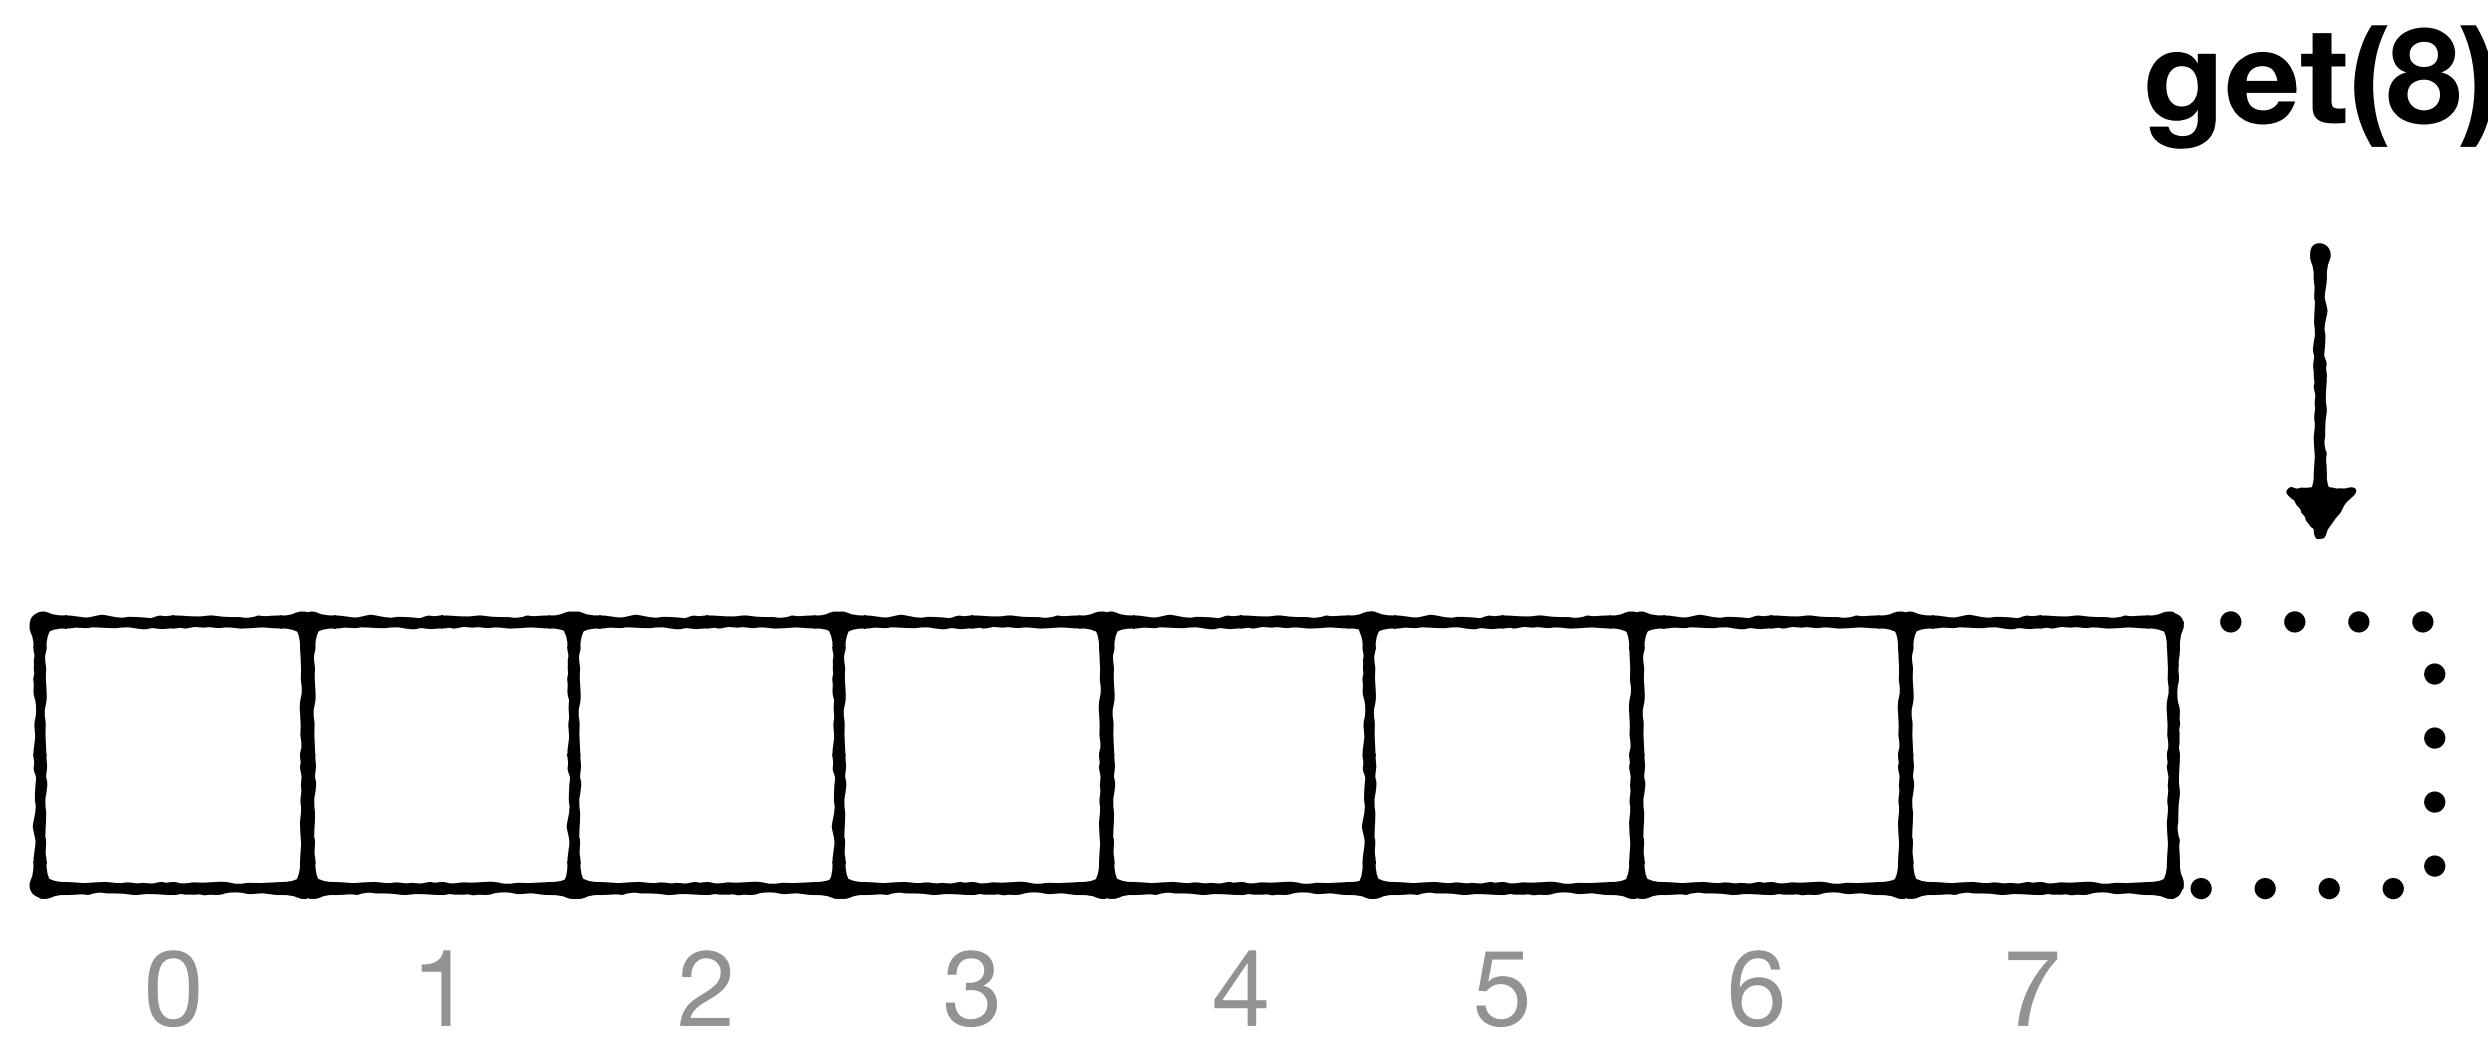


Supports random access

get(8)



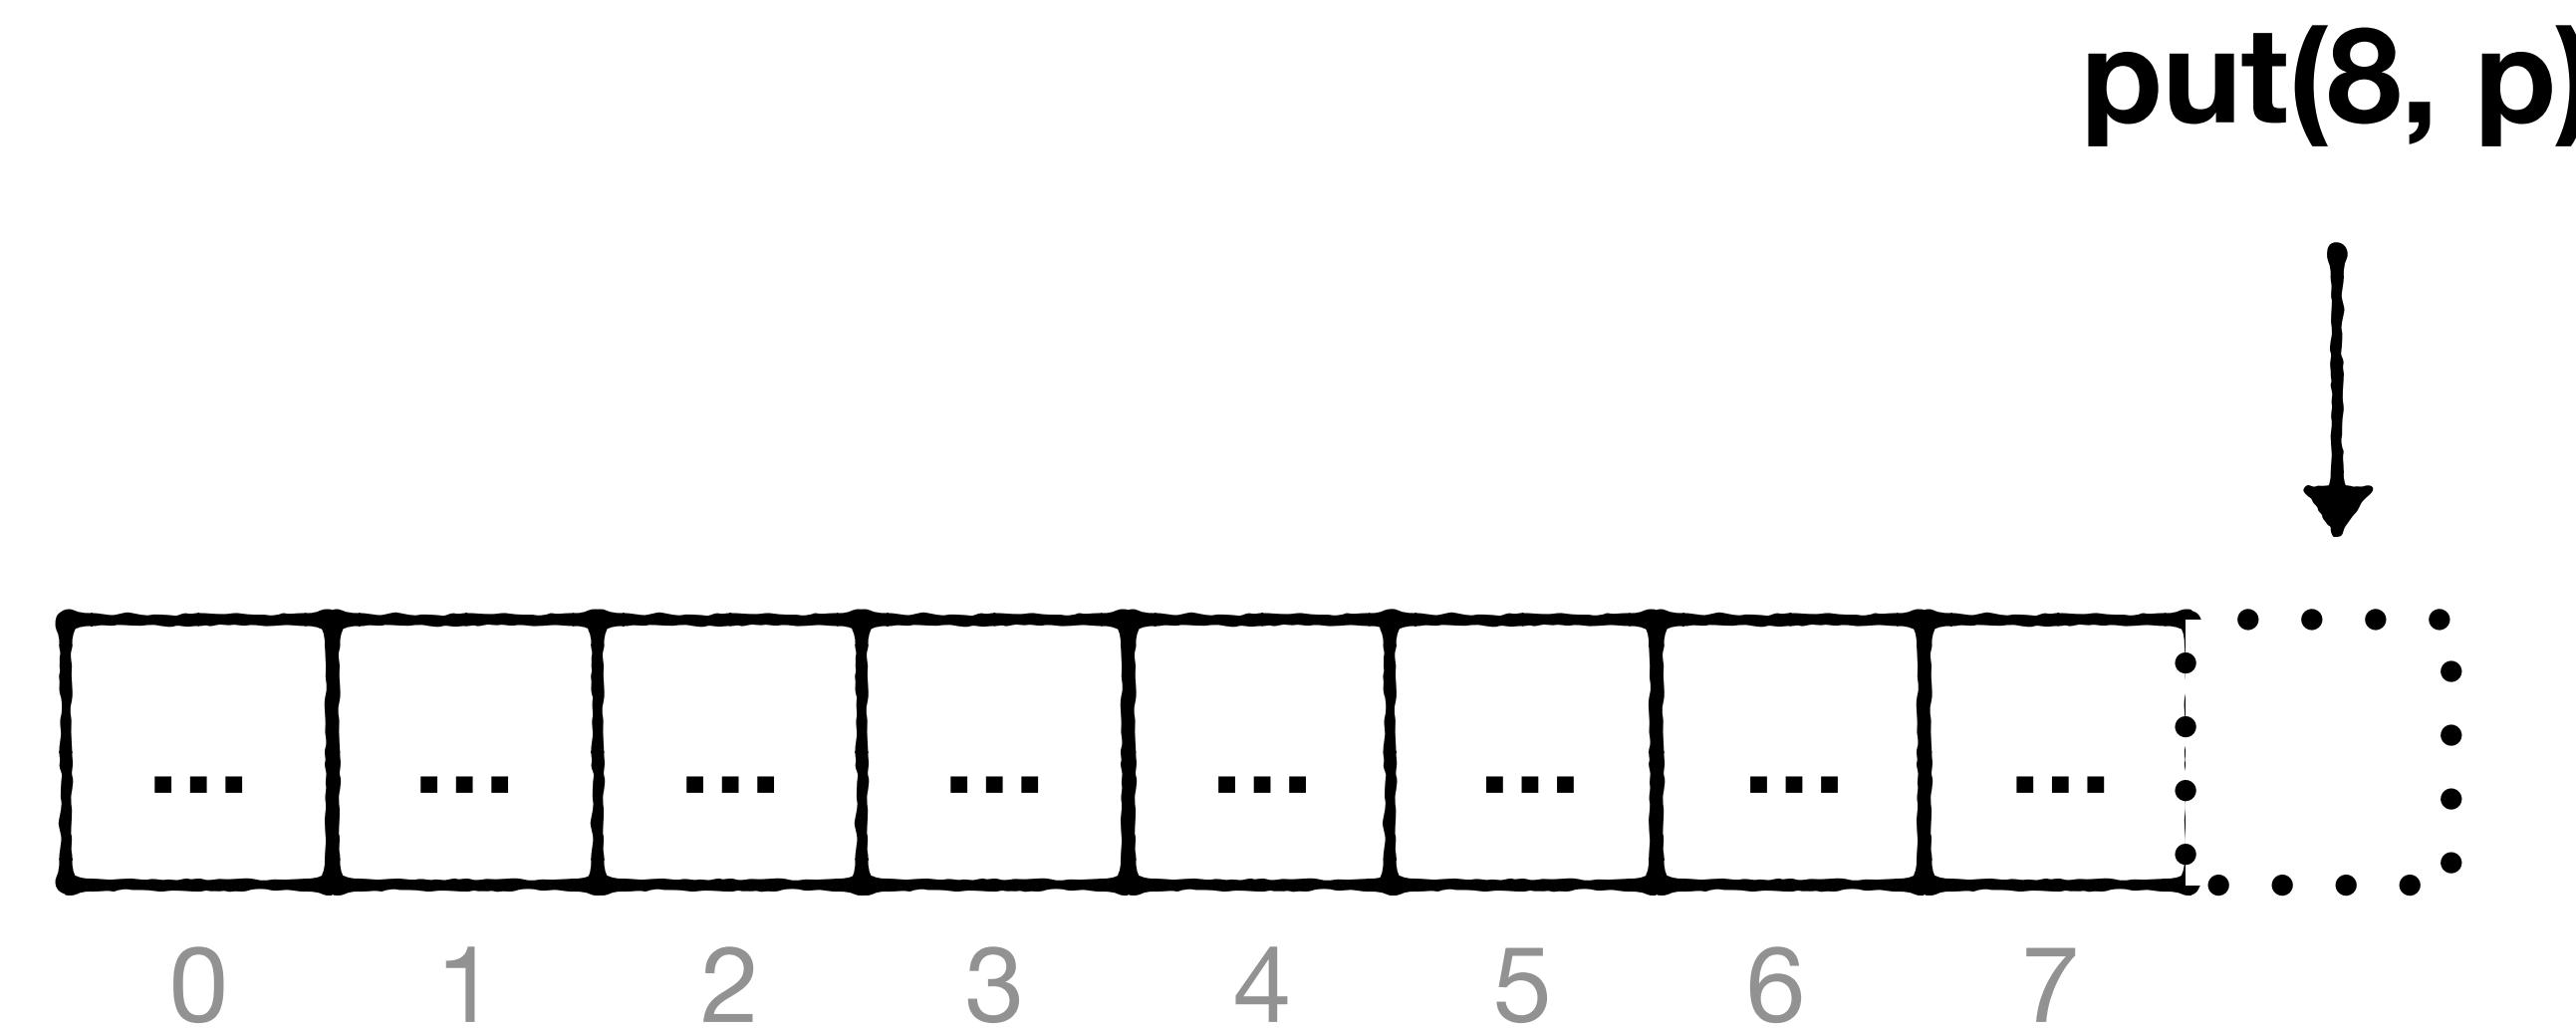
Overflow error (e.g., java)



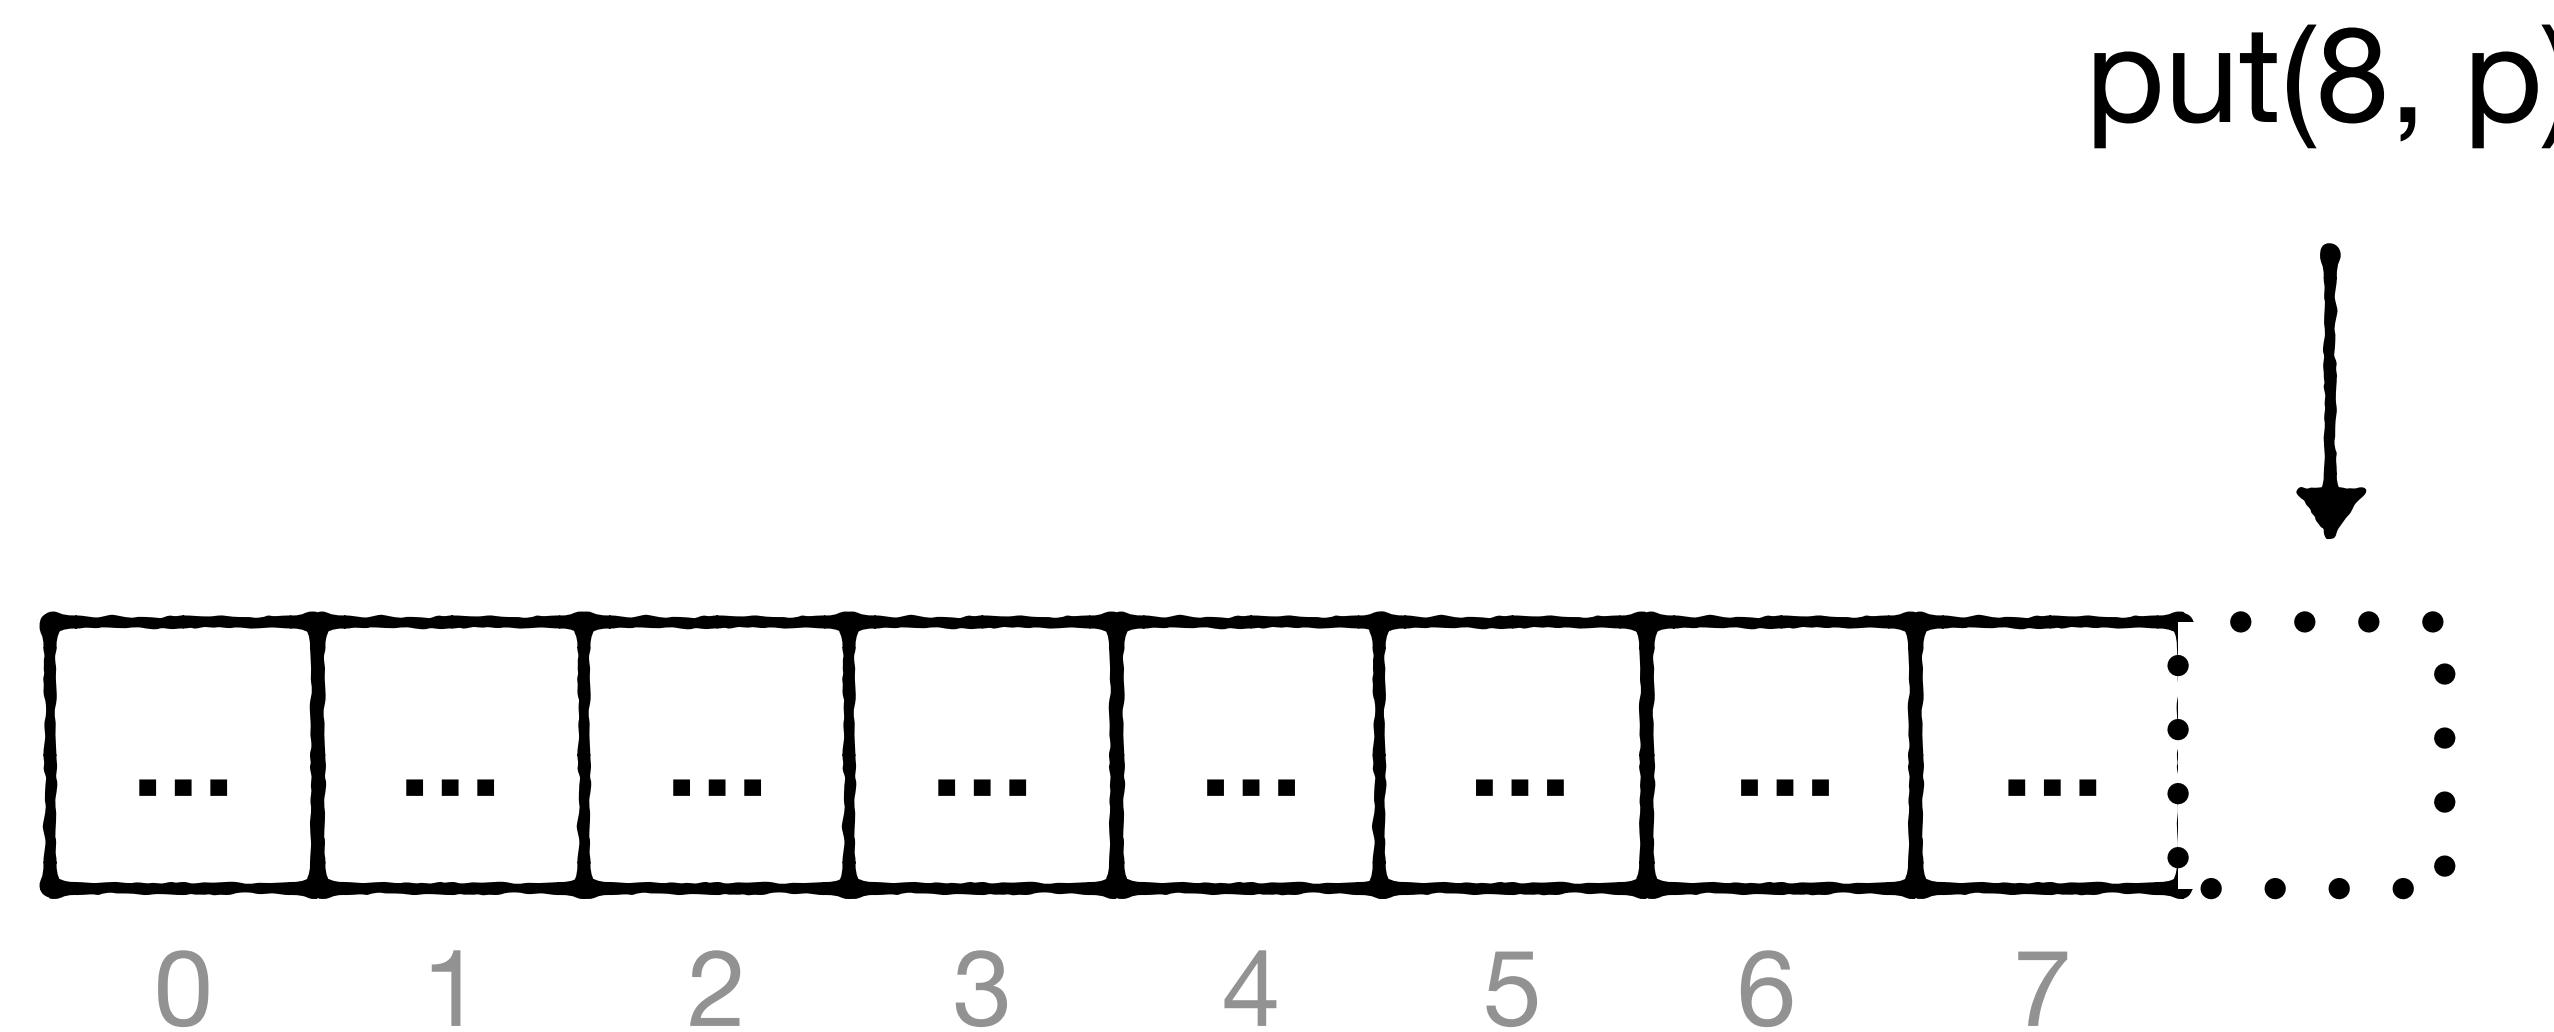
Overflow error (e.g., java)

Undefined behavior (e.g., C++)

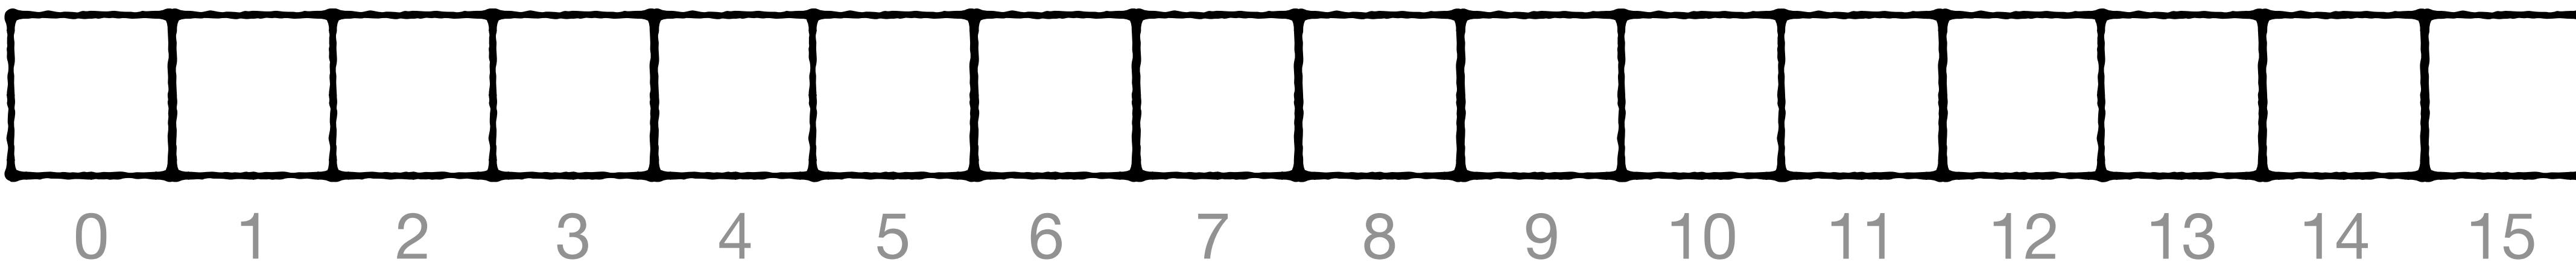
How to keep inserting when out of space?

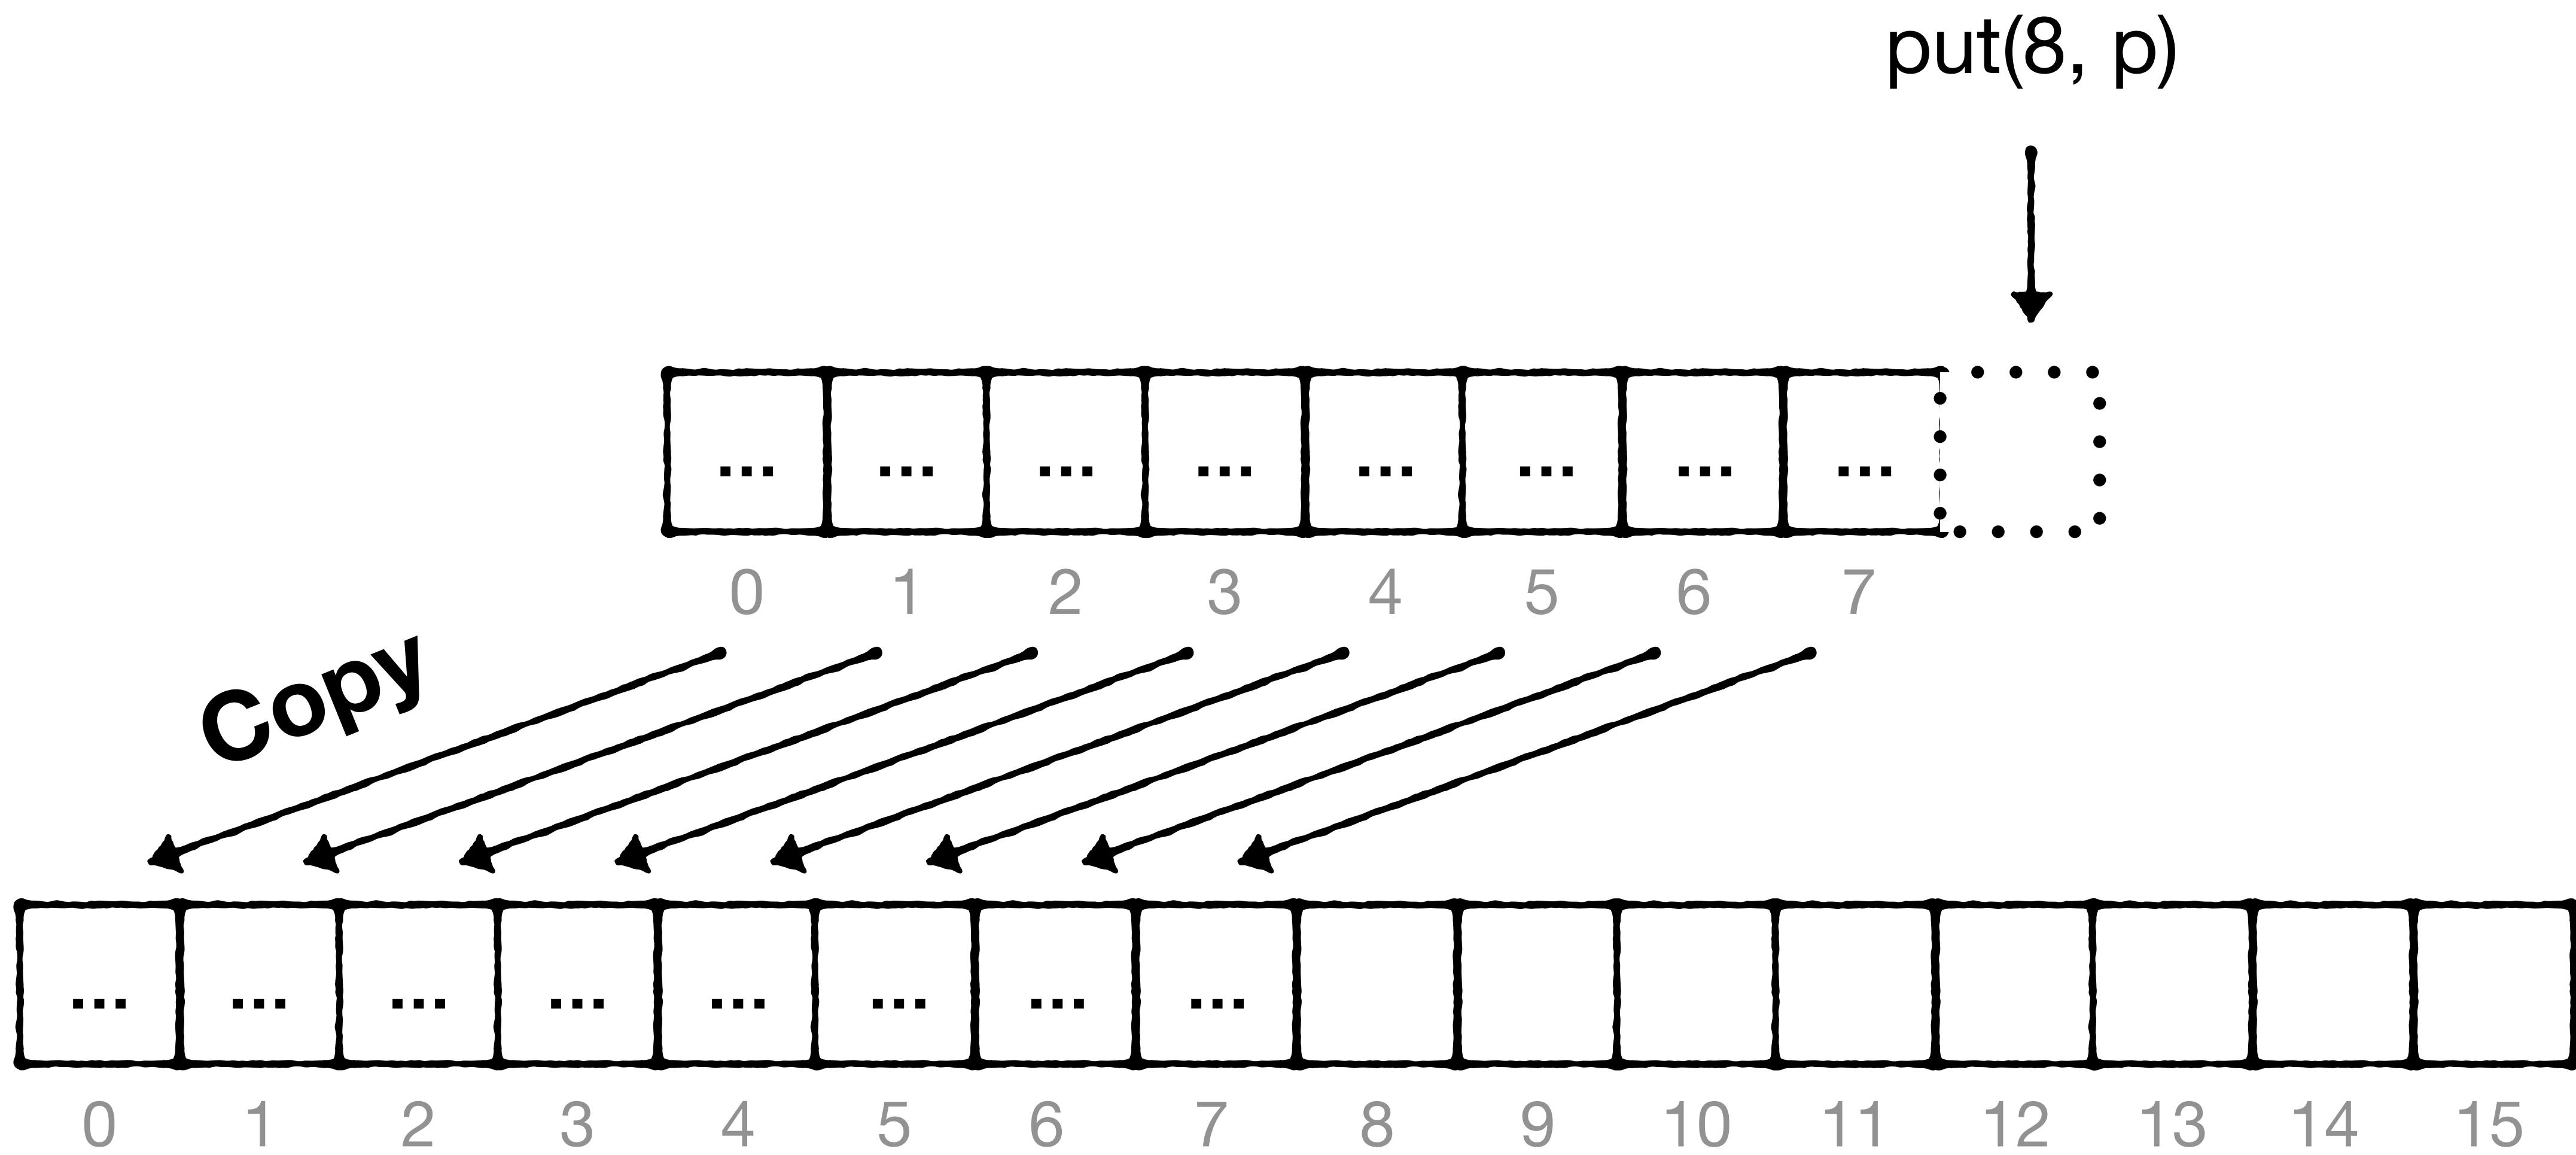


How to keep inserting when out of space?

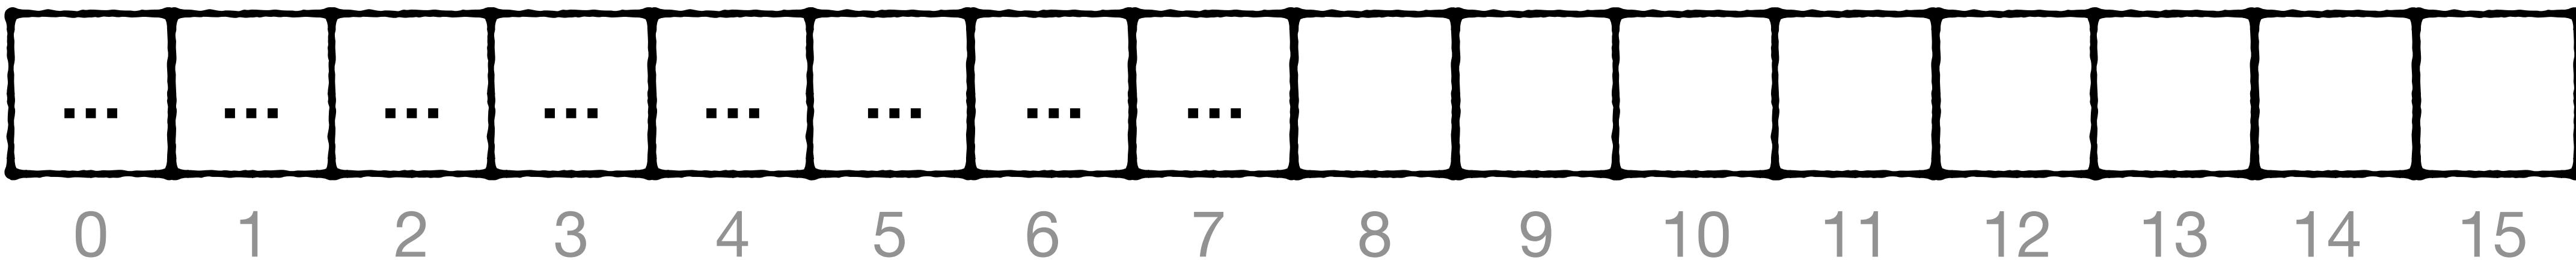
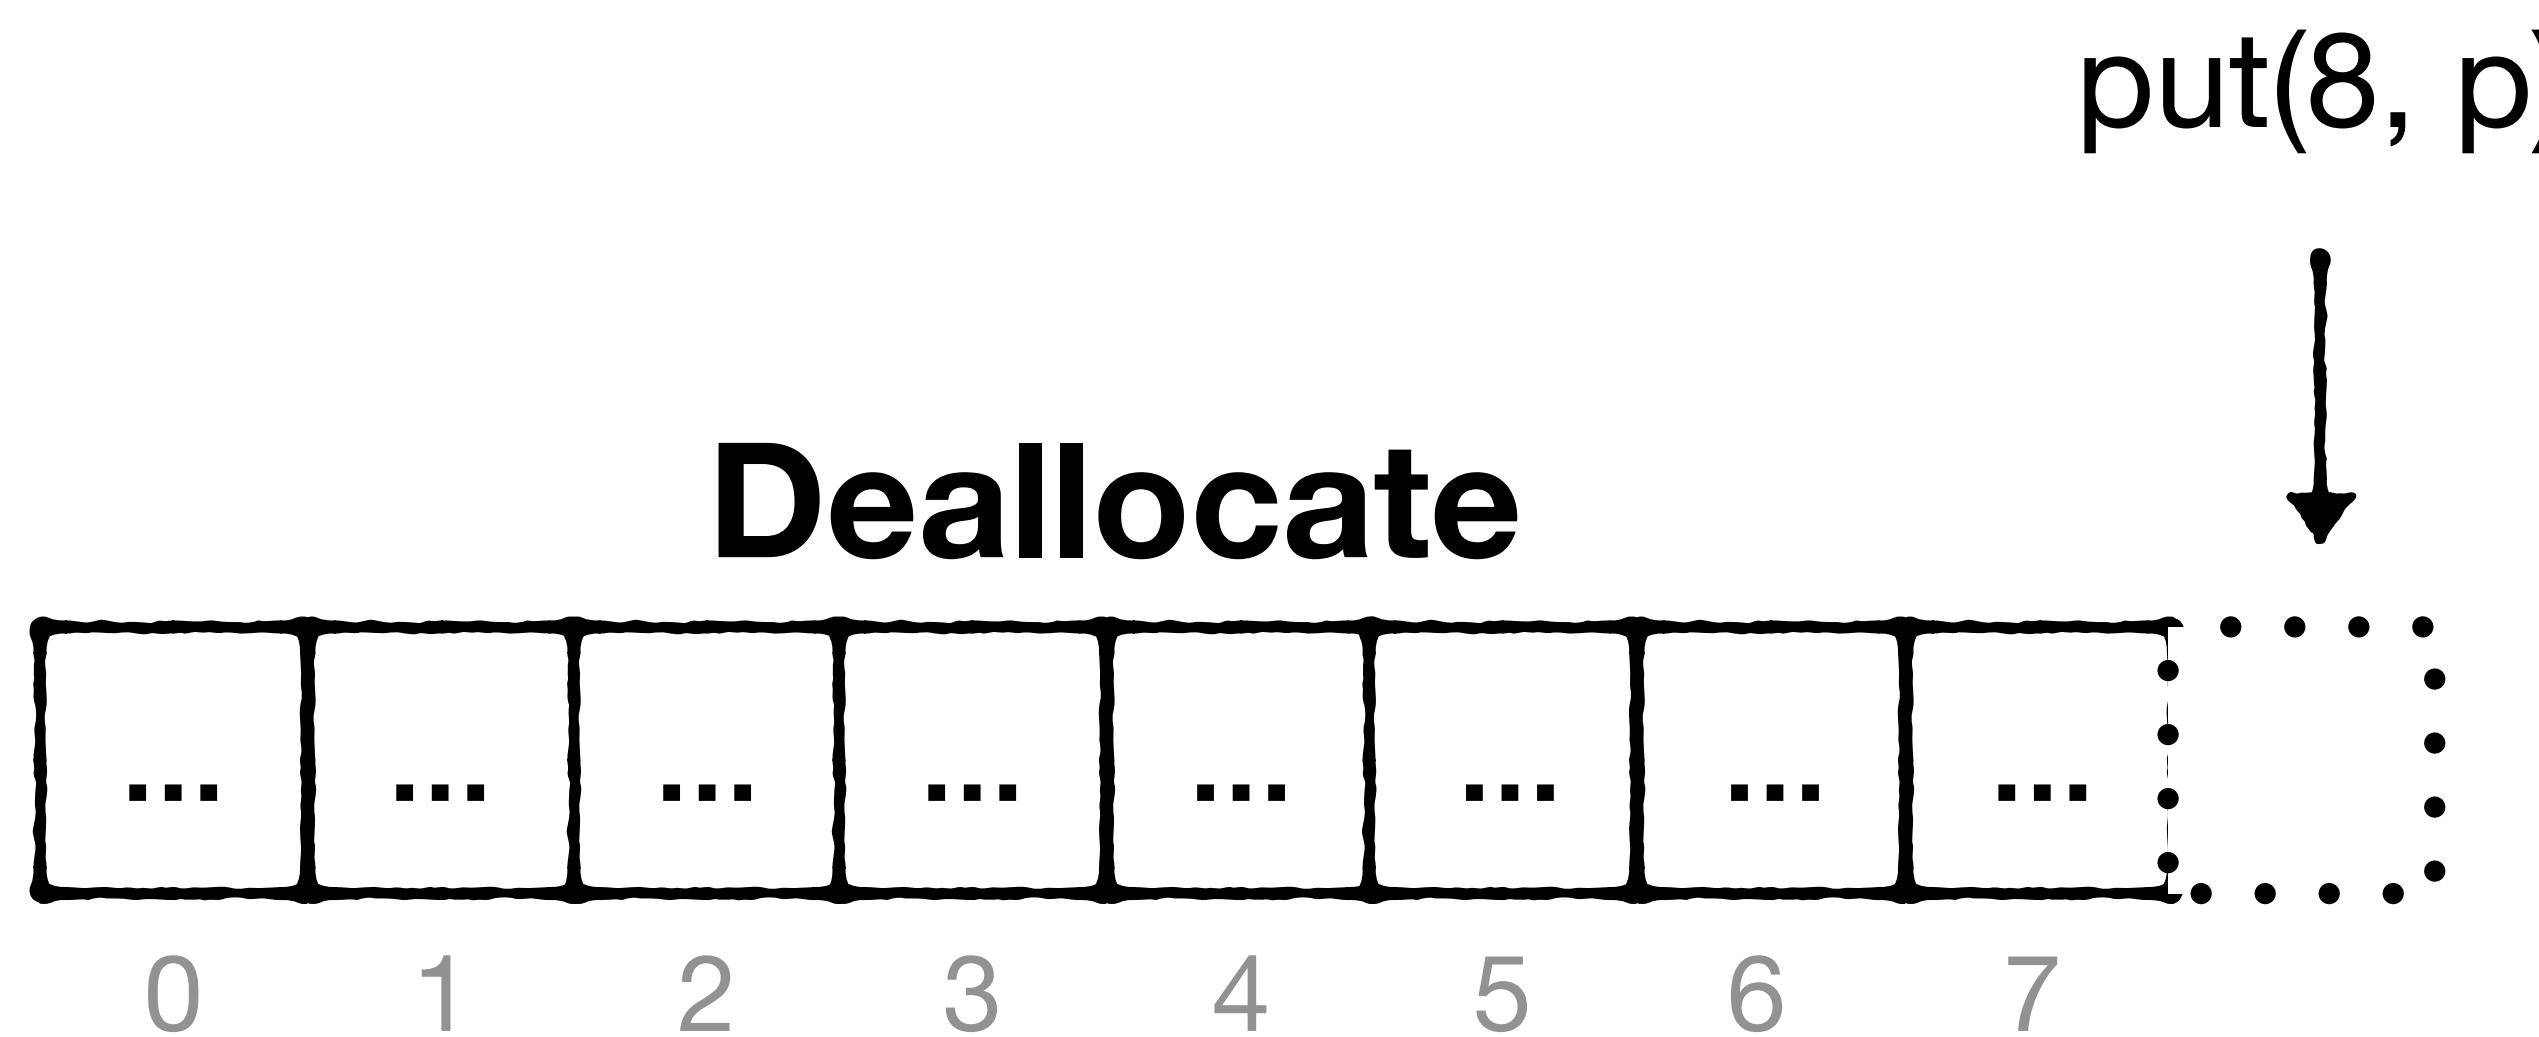


Allocate

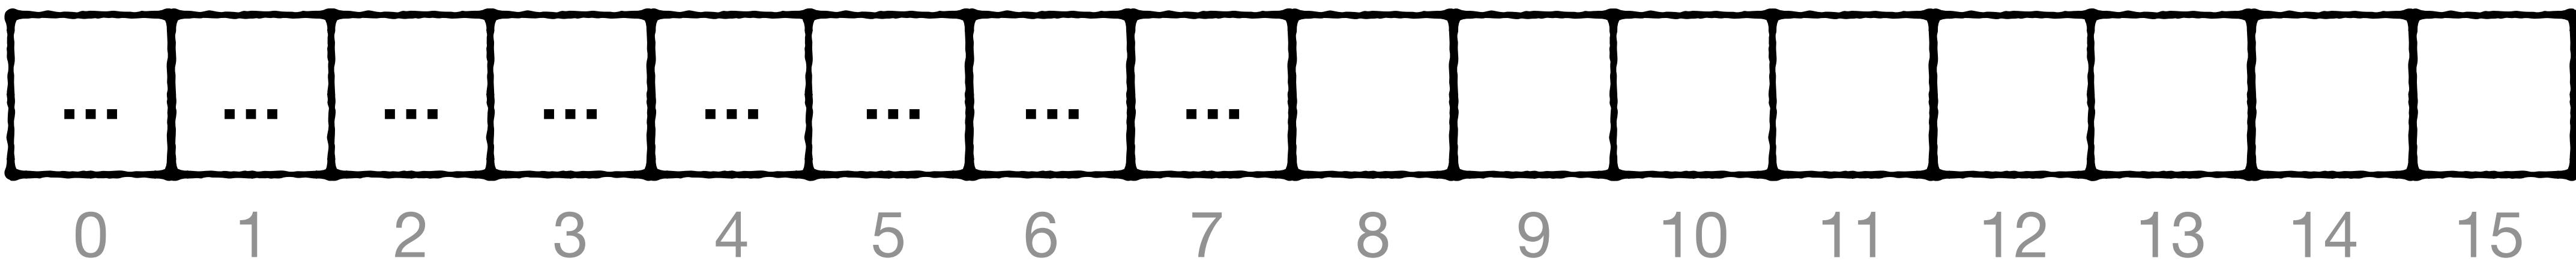




Deallocate



put(8, p)



C++ offers static & dynamic arrays

Static

```
int nums[4] = {0, 0 ,0 ,0};
```

Dynamic

```
std::vector<int> vec;
```

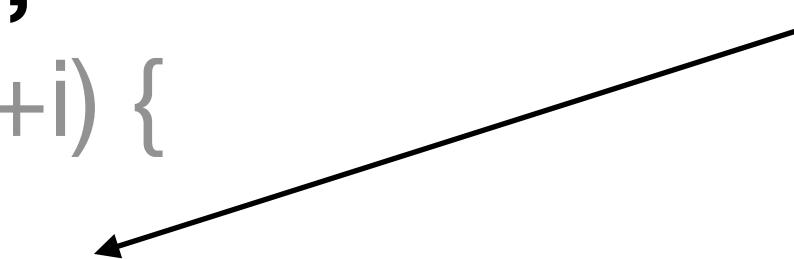
C++ offers static & dynamic arrays

```
std::vector<int> vec;
for (int i = 0; i < 10; ++i) {
    vec.push_back(i);
    std::cout << "Added " << i << ", size: " << vec.size()
        << ", capacity: " << vec.capacity() << std::end;
}
```

C++ offers static & dynamic arrays

```
std::vector<int> vec;  
for (int i = 0; i < 10; ++i) {  
    vec.push_back(i);  
    std::cout << "Added " << i << ", size: " << vec.size()  
        << ", capacity: " << vec.capacity() << std::end;  
}
```

Resize if we exceed capacity



C++ offers static & dynamic arrays

```
std::vector<int> vec;
for (int i = 0; i < 10; ++i) {
    vec.push_back(i);
    std::cout << "Added " << i << ", size: " << vec.size()
        << ", capacity: " << vec.capacity() << std::endl;
}
```

Element	Size	Capacity
0	1	1
1	2	2
2	3	4
3	4	4
4	5	8
5	6	8
6	7	8
7	8	8
8	9	16
9	10	16

C++ offers static & dynamic arrays

```
std::vector<int> vec;
for (int i = 0; i < 10; ++i) {
    vec.push_back(i);
    std::cout << "Added " << i << ", size: " << vec.size()
        << ", capacity: " << vec.capacity() << std::endl;
}
```

Element	Size	Capacity
0	1	1
1	2	2
2	3	4
3	4	4
4	5	8
5	6	8
6	7	8
7	8	8
8	9	16
9	10	16

Uses Growth Factor 2

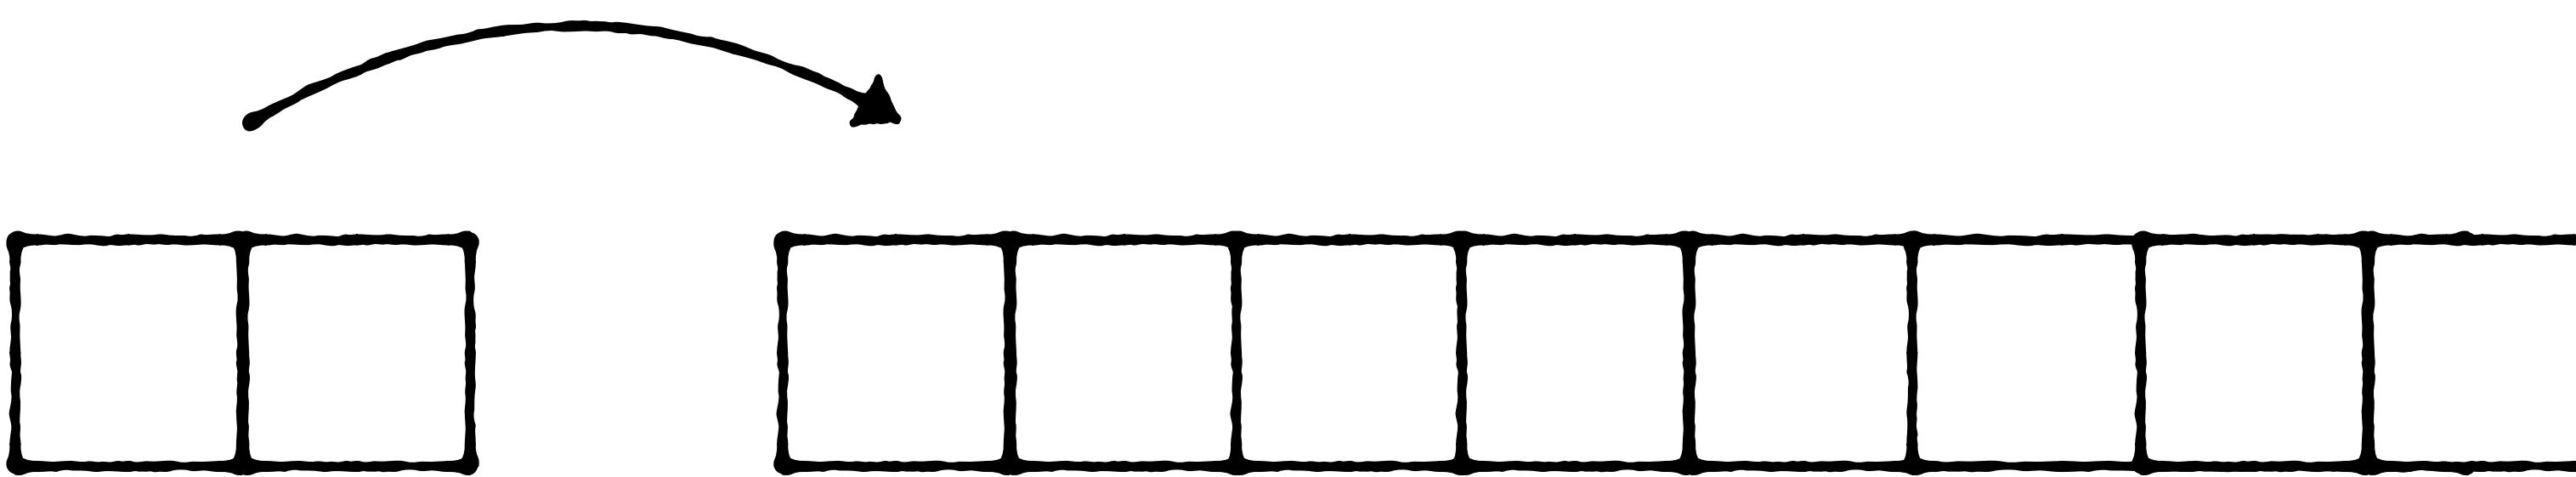
What if Growth Factor G is

too high?

too low?

What if Growth Factor G is **too high**

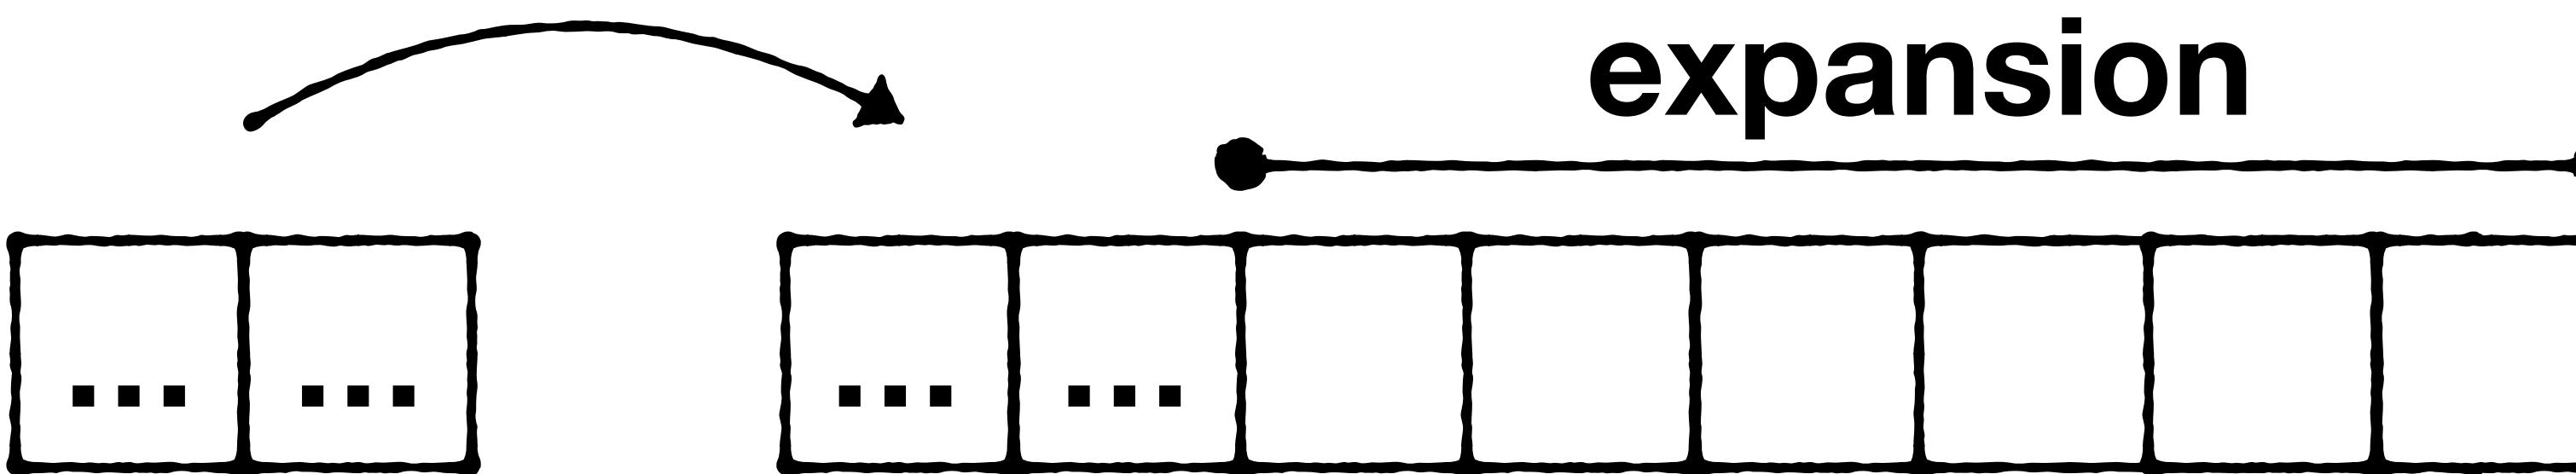
e.g., 4



What if Growth Factor G is **too high**

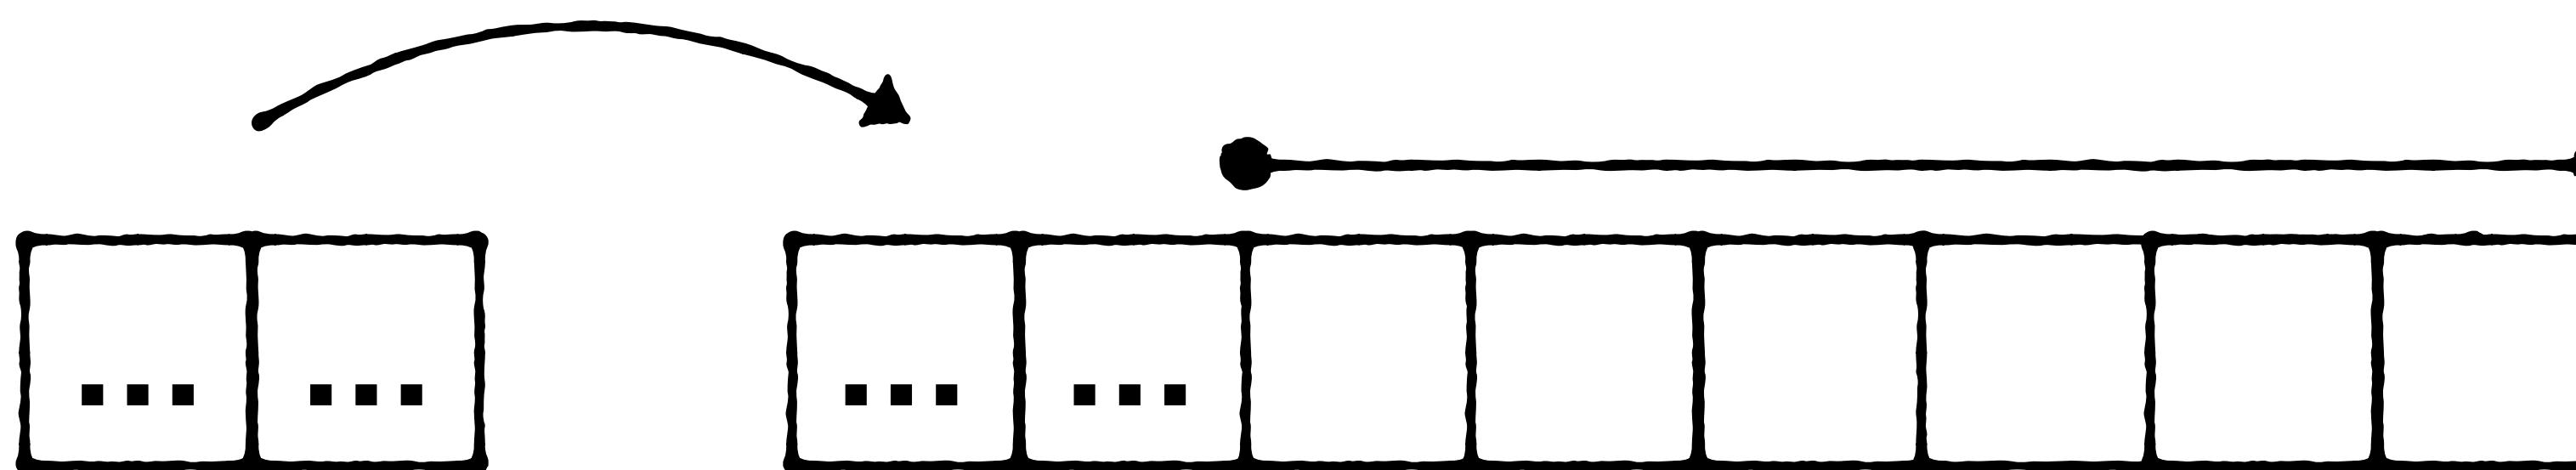
e.g., 4

**x4 space wasted after
expansion**



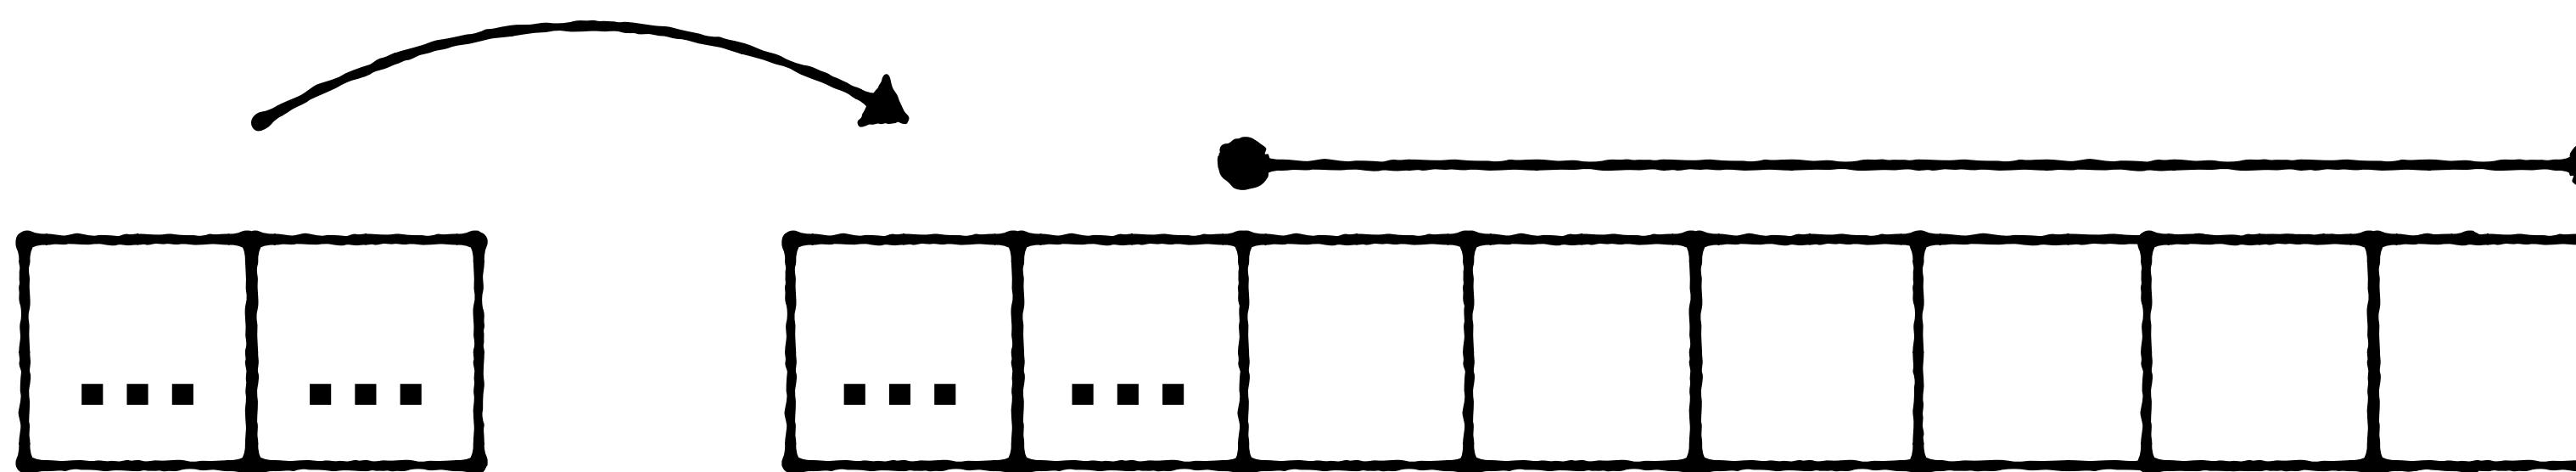
What if Growth Factor G is **too high**

Space-amplification =
$$\frac{\text{Physical space used}}{\text{Data size}}$$



What if Growth Factor G is too high

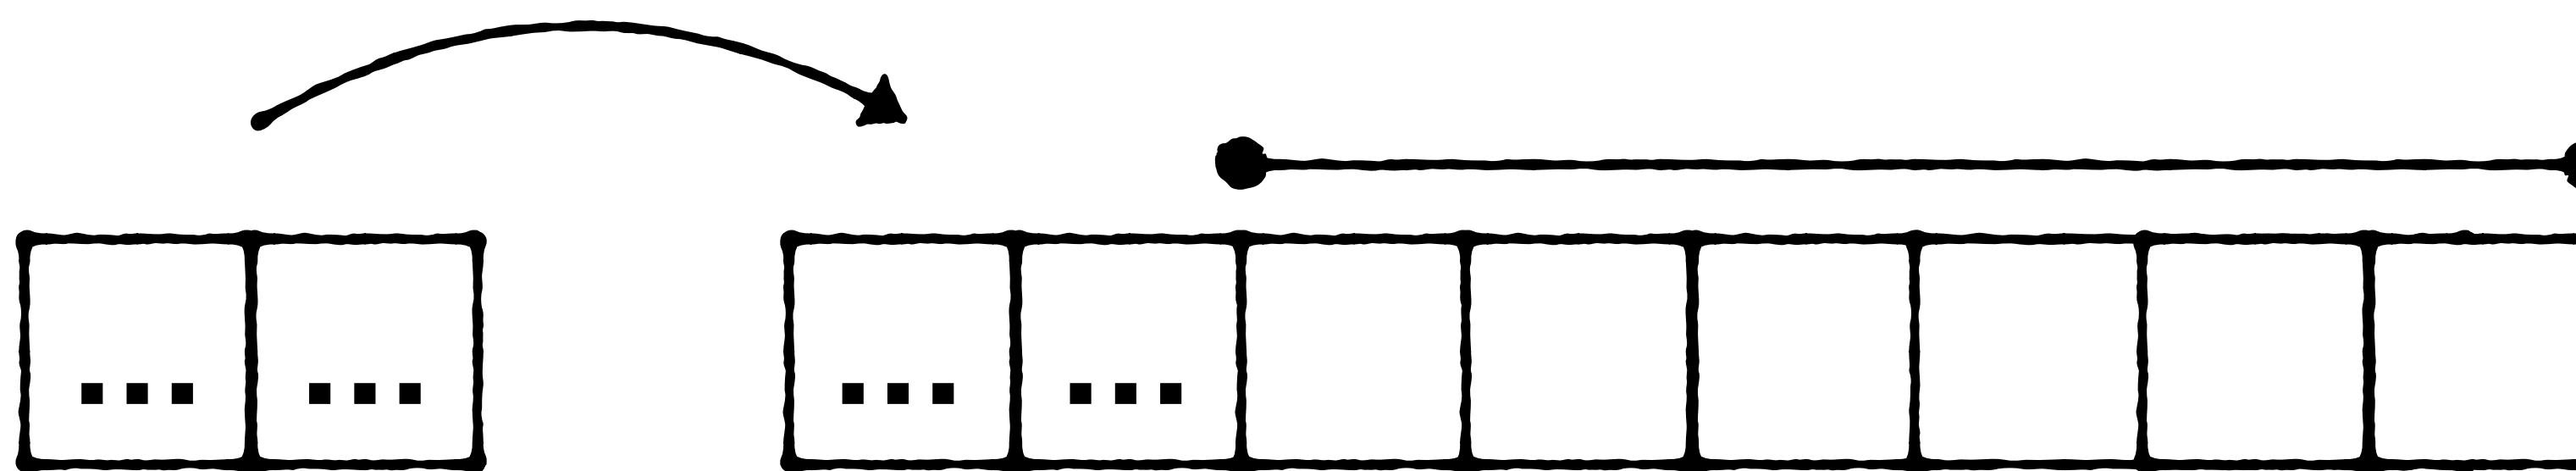
$$\text{Space-amplification} = \frac{\text{Physical space used}}{\text{Data size}} = G$$



What if Growth Factor G is **too high**

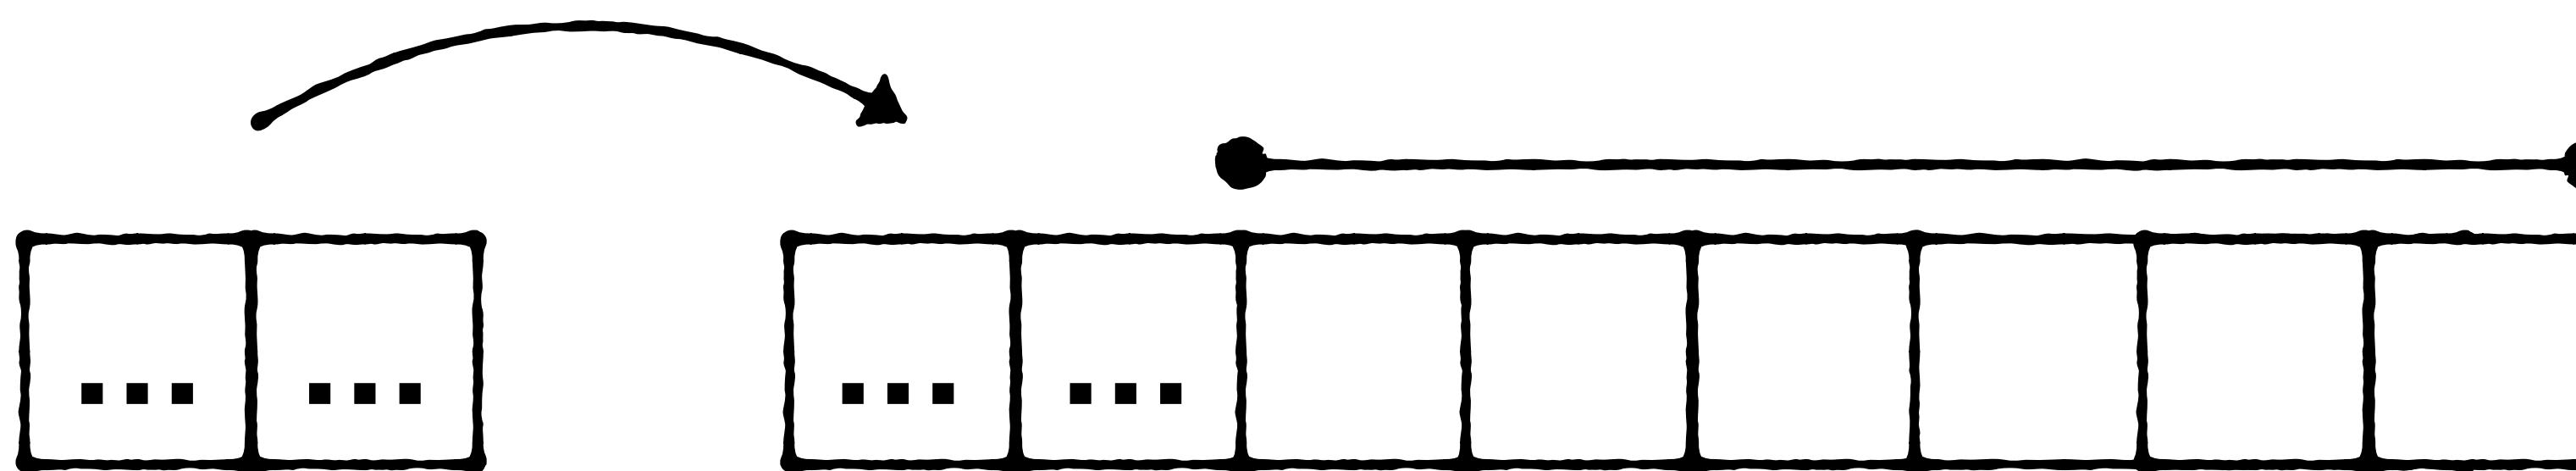
Space-amplification = **G**

**Right after
expansion**



What if Growth Factor G is too high

**Max
Space-amplification** = $G + 1$ **During
expansion**

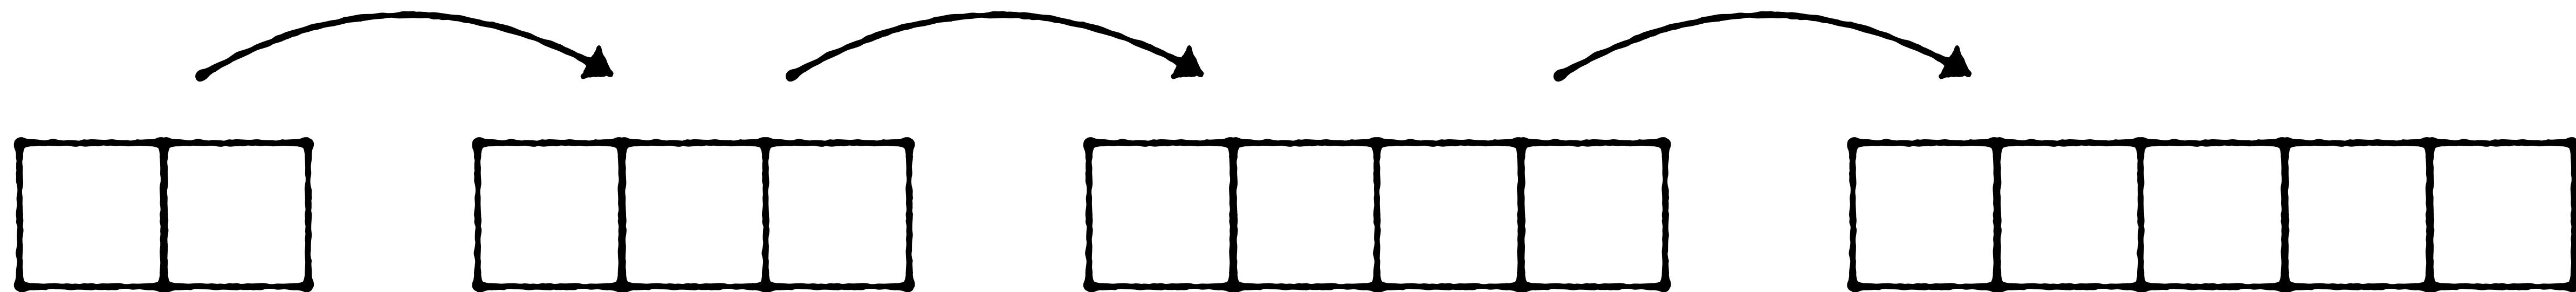


What if Growth Factor G is **too low**

e.g., 1.2

What if Growth Factor G is **too low**

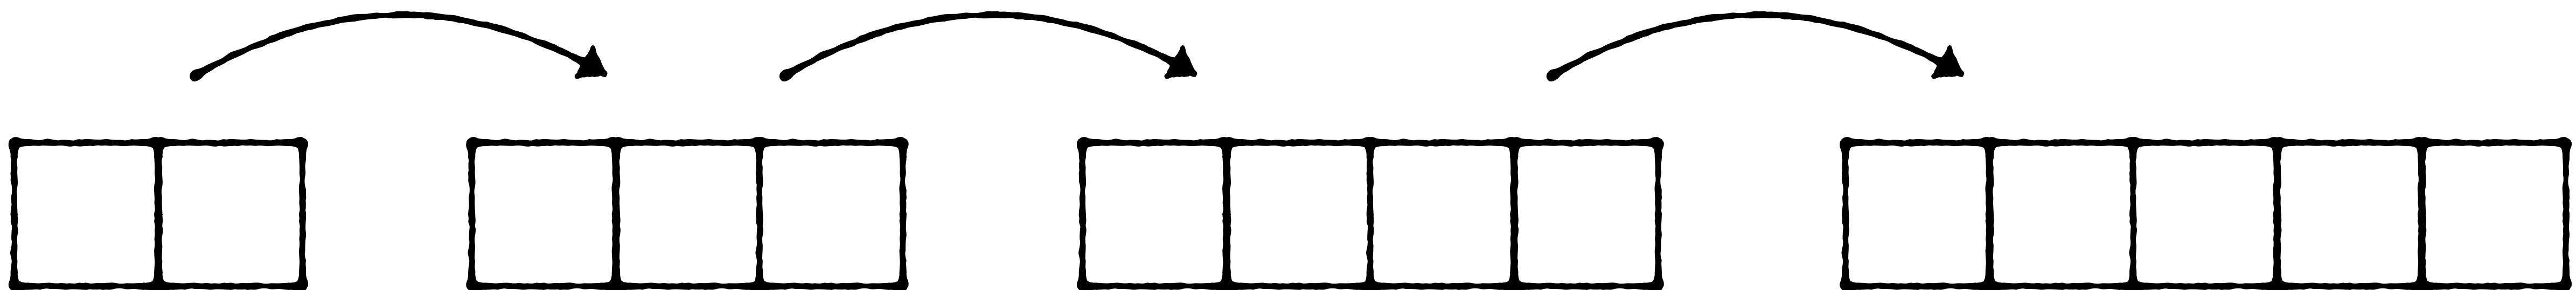
e.g., 1.2



What if Growth Factor G is **too low**

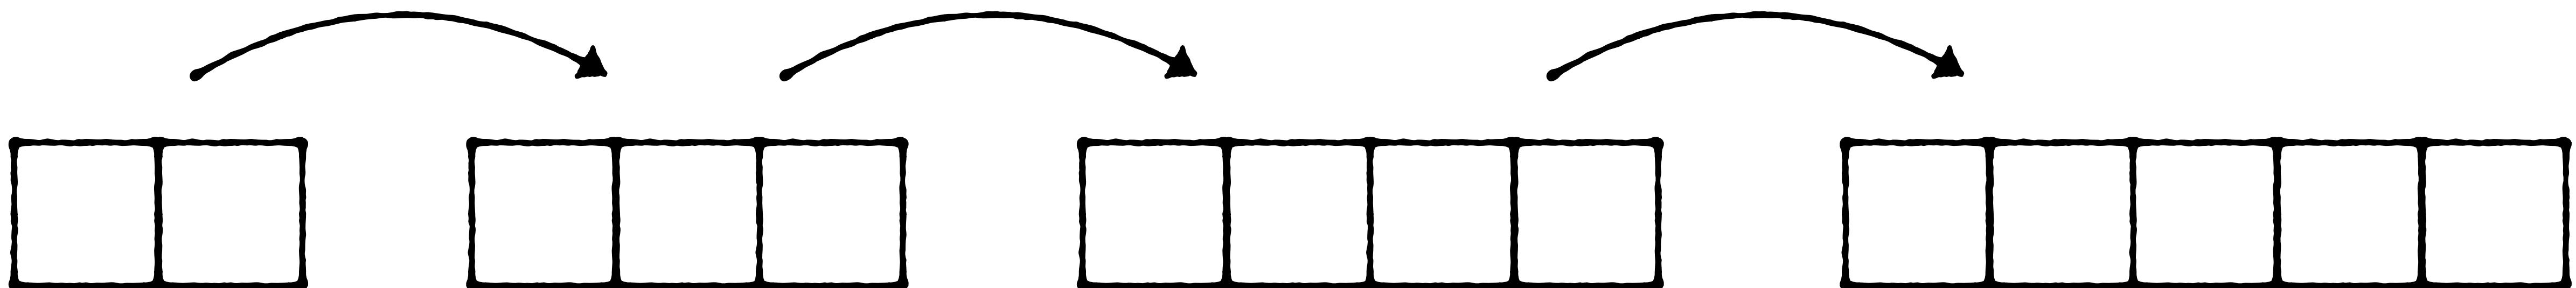
e.g., 1.2

Insertion overheads increase



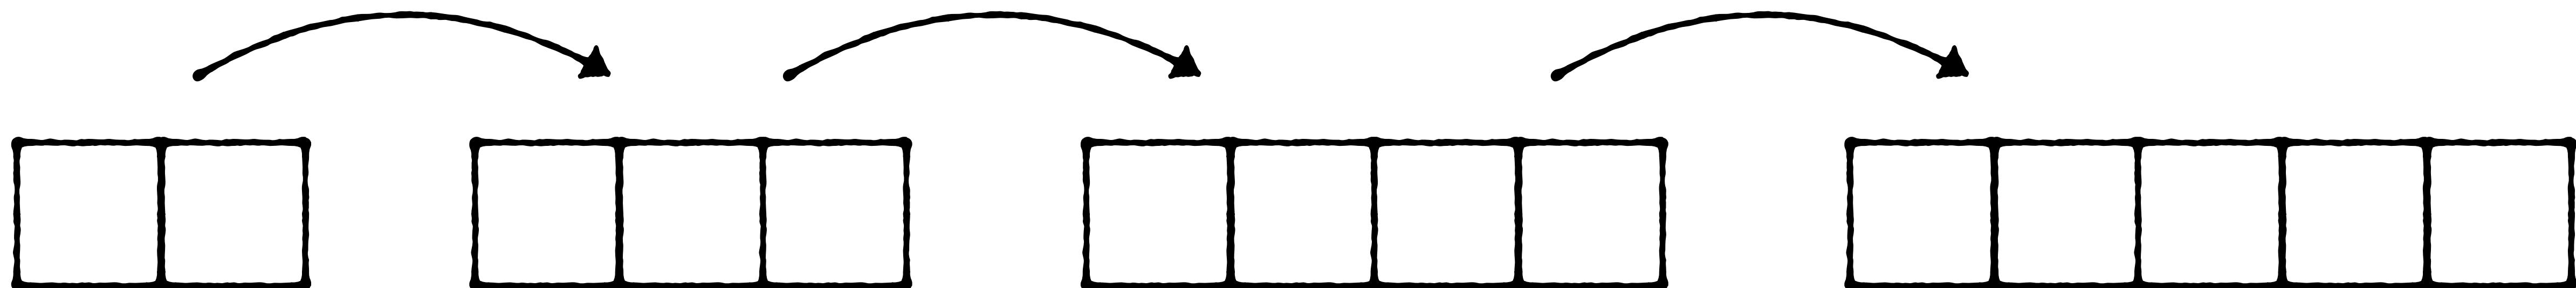
What if Growth Factor G is **too low**

$$\text{Write-amplification} = \frac{\text{Physical data written}}{\text{Data size}}$$



What if Growth Factor G is **too low**

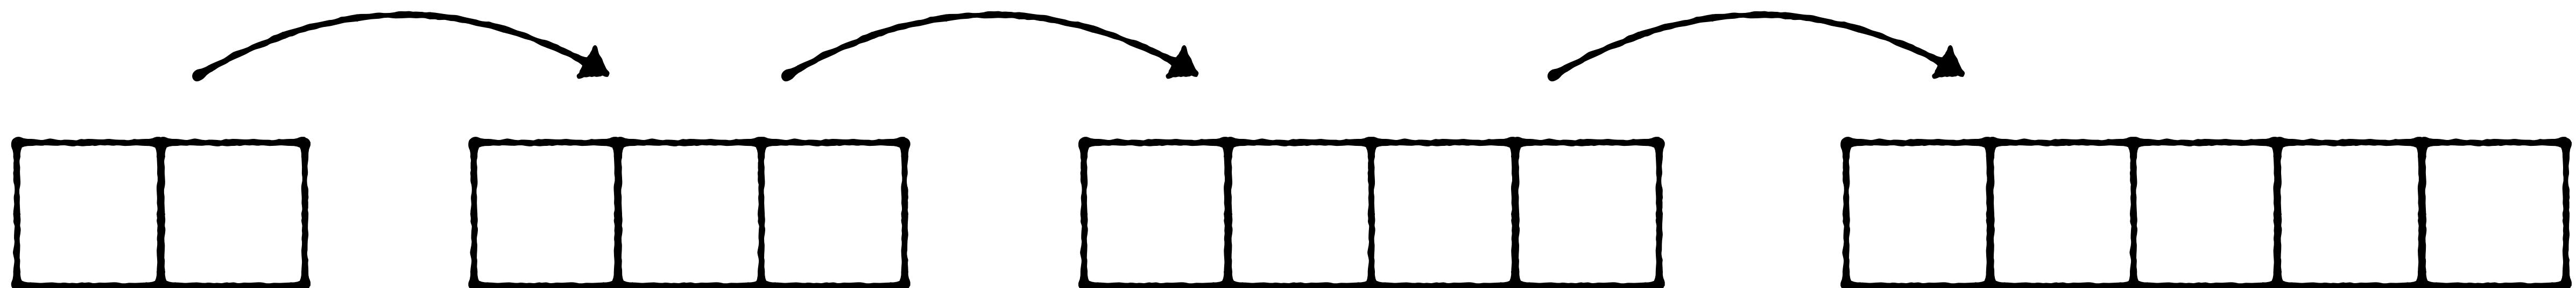
$$\text{Write-amplification} = \frac{\text{Physical data written}}{\text{Data size}} = \frac{G}{G - 1}$$



What if Growth Factor G is **too low**

$$\text{Write-amplification} = \frac{\text{Physical data written}}{\text{Data size}} = \frac{G}{G - 1}$$

Geometric series sum



Growth factor G impact

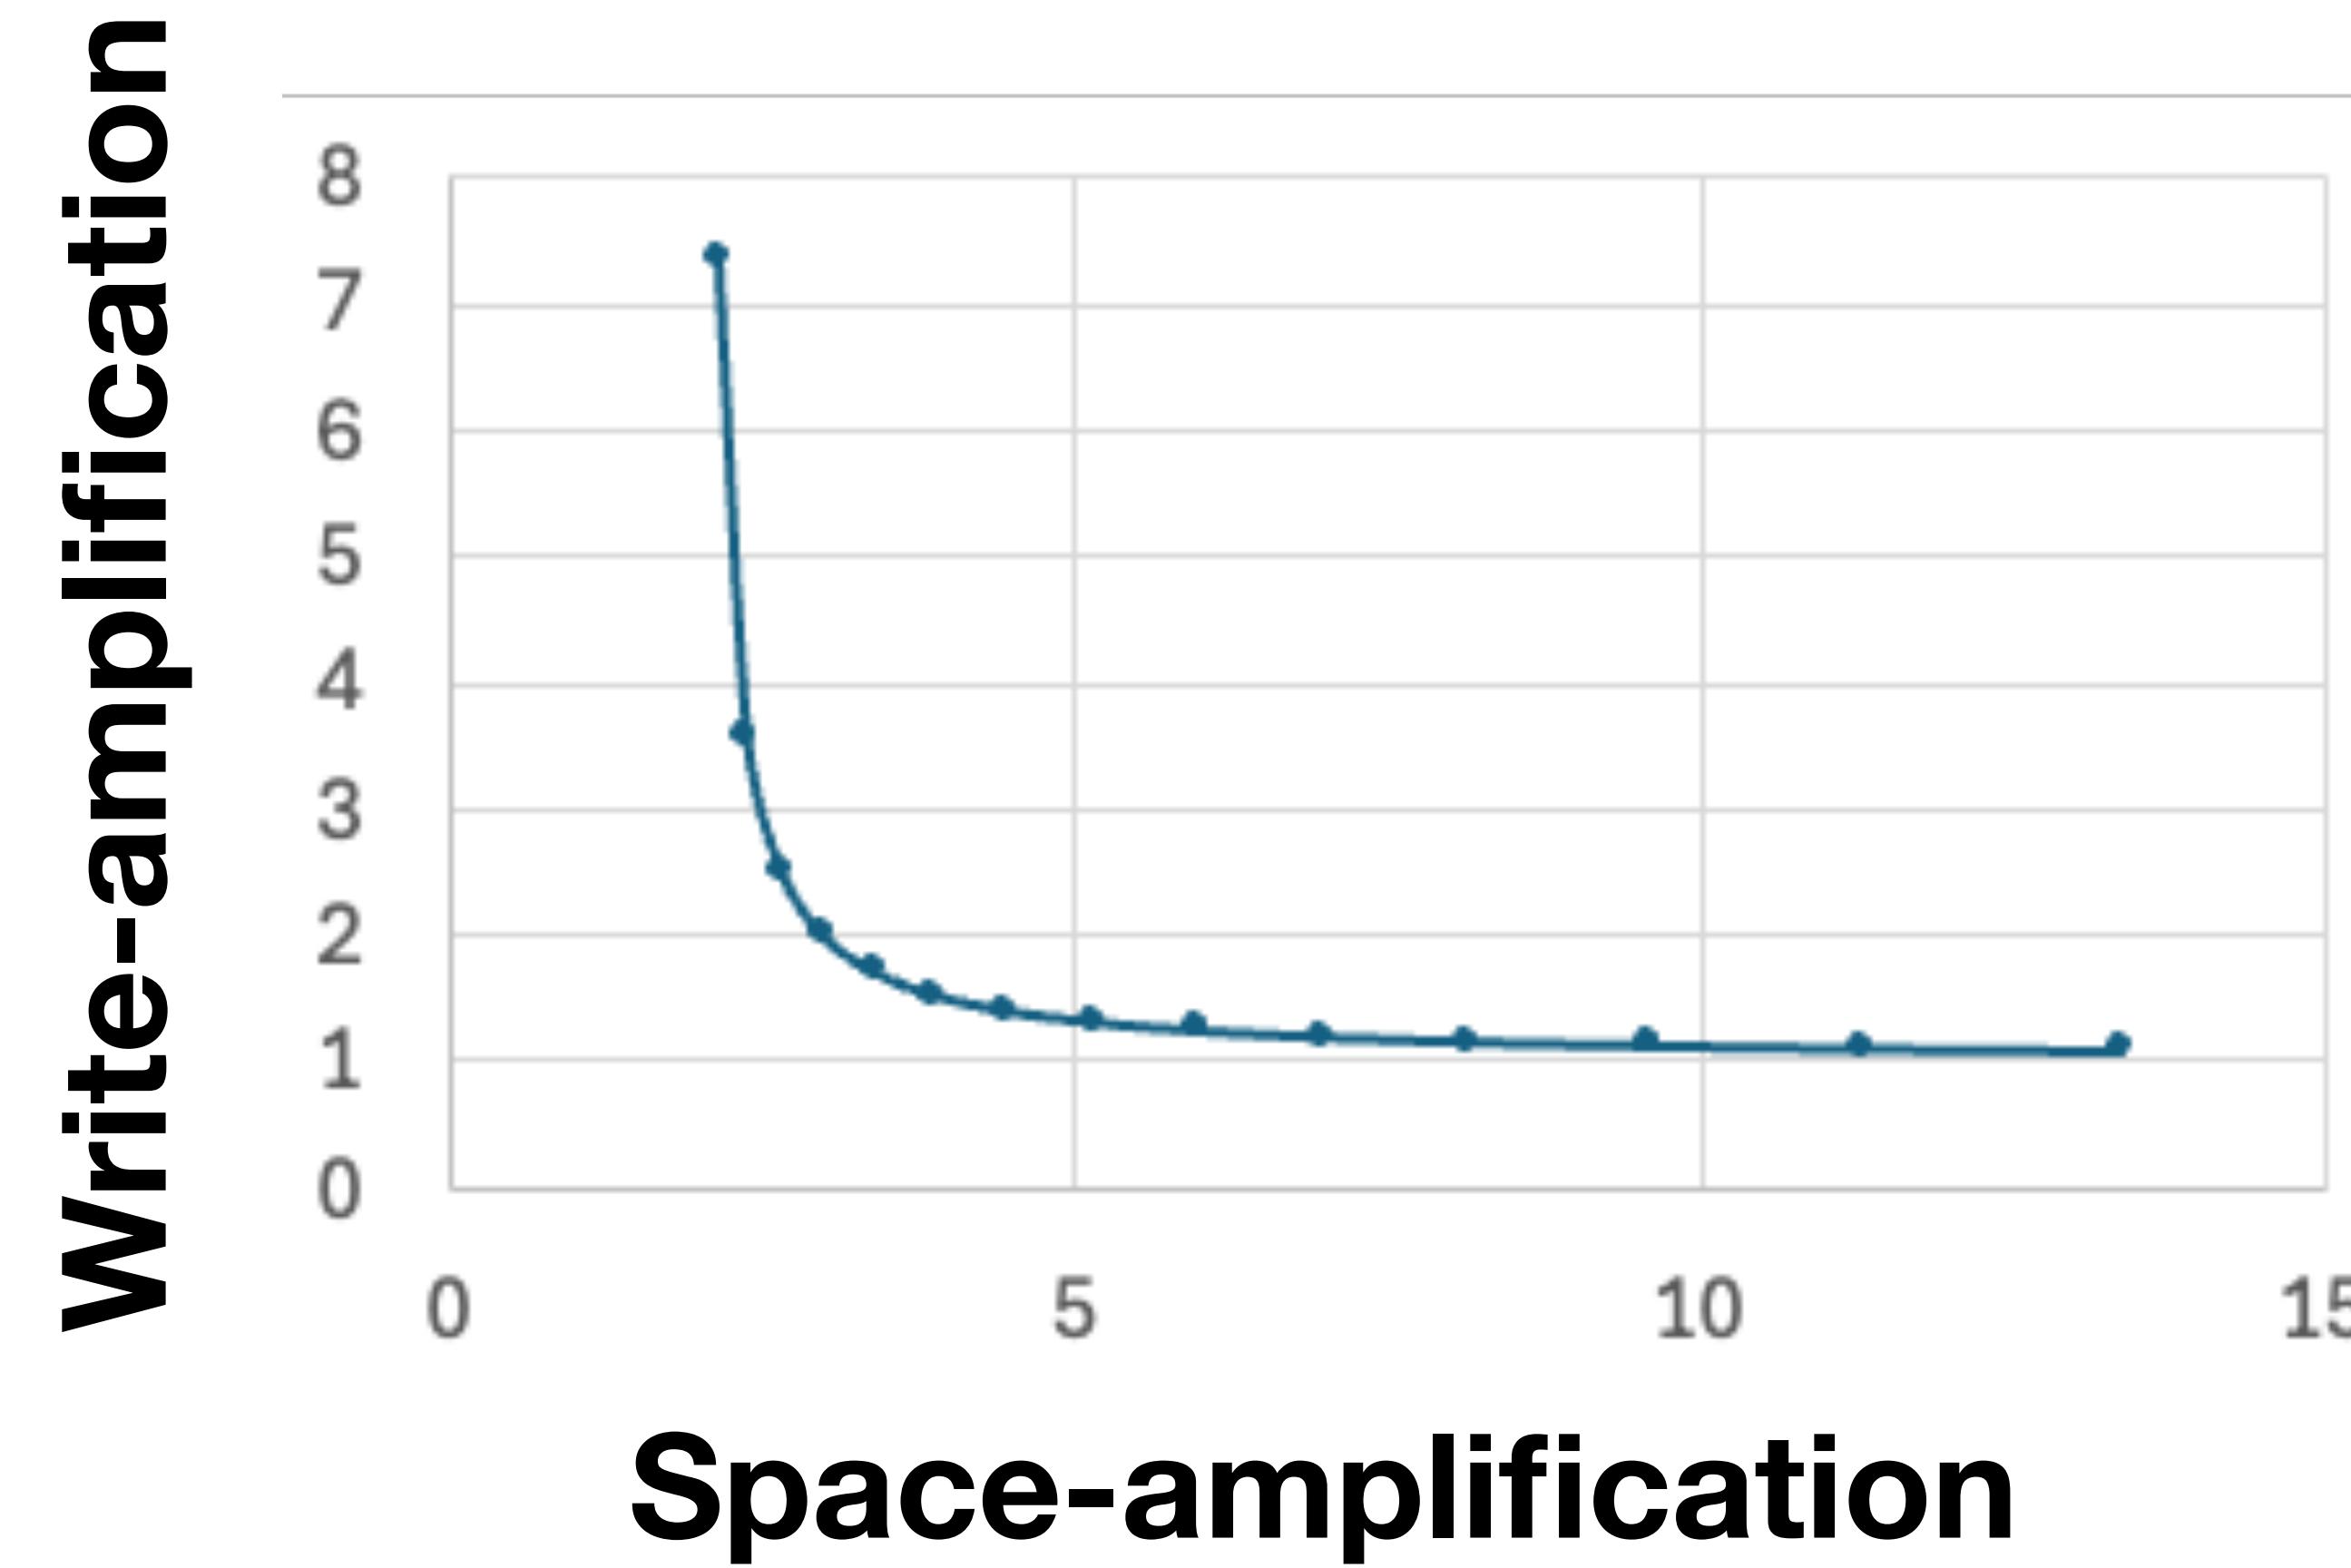
Space-amplification

$G+1$

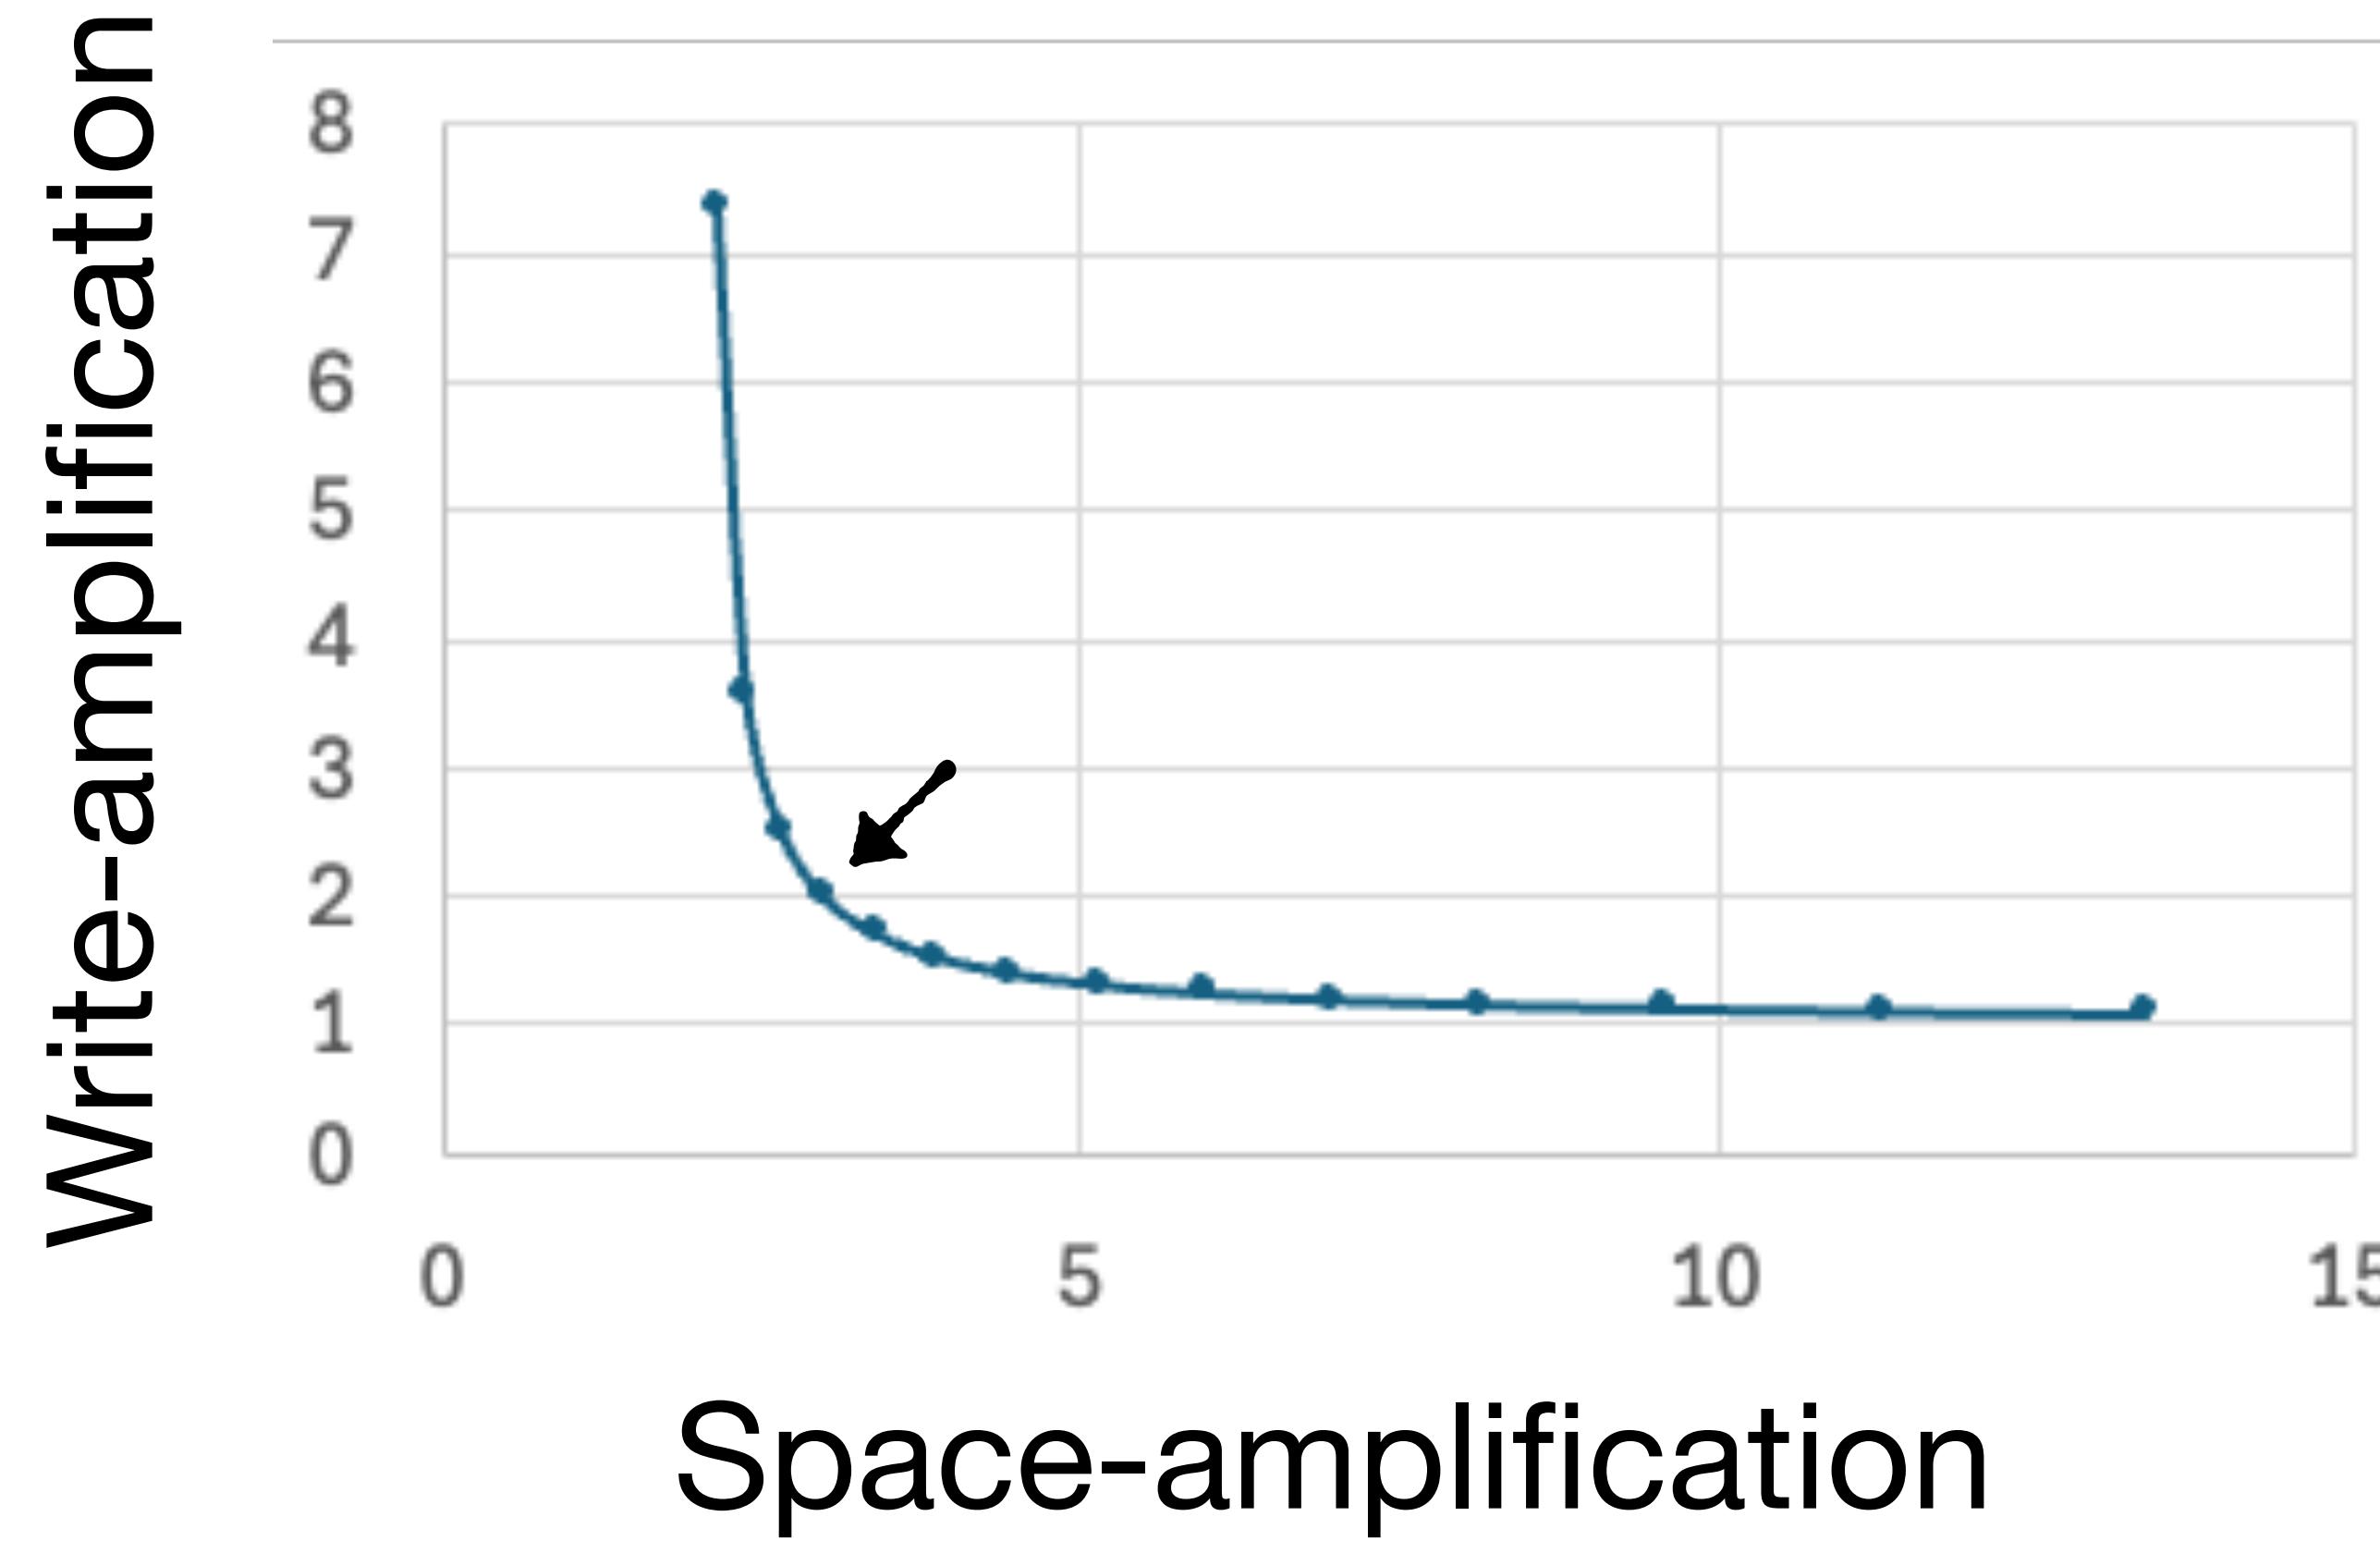
Write-amplification

$\frac{G}{G - 1}$

Growth factor G impact

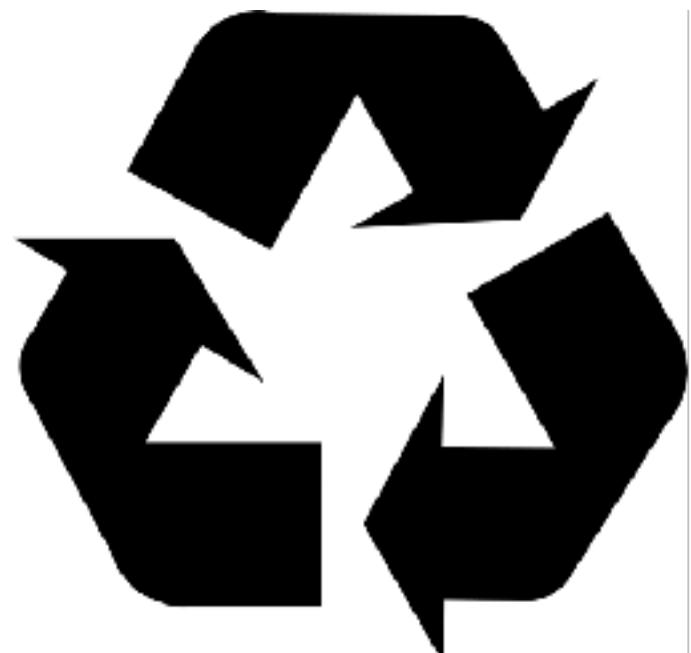


Growth factor 2 achieves a good balance

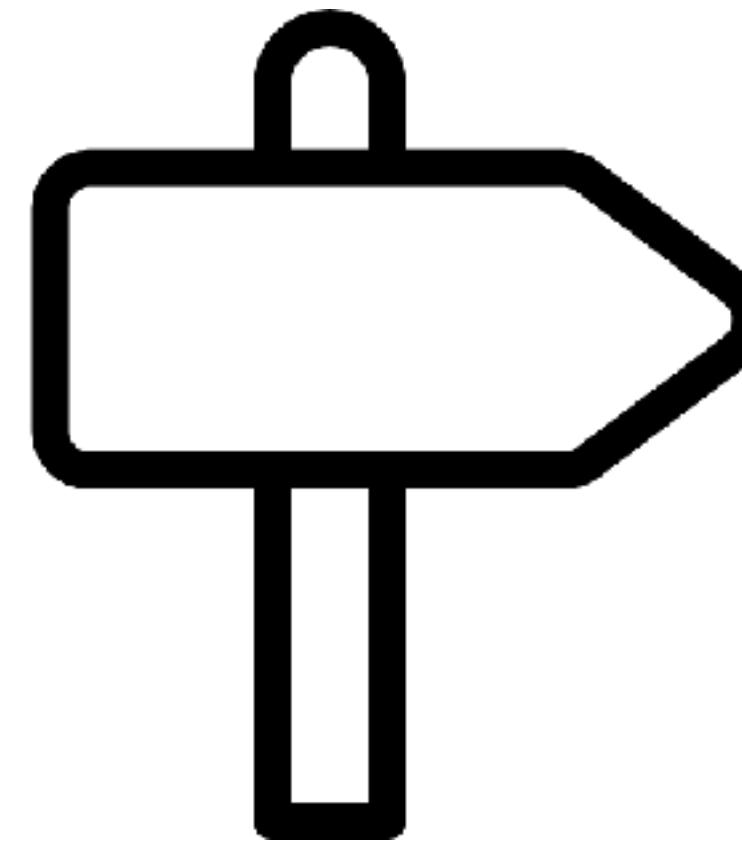


And now to new stuff

**Reusing Deallocated
Space**

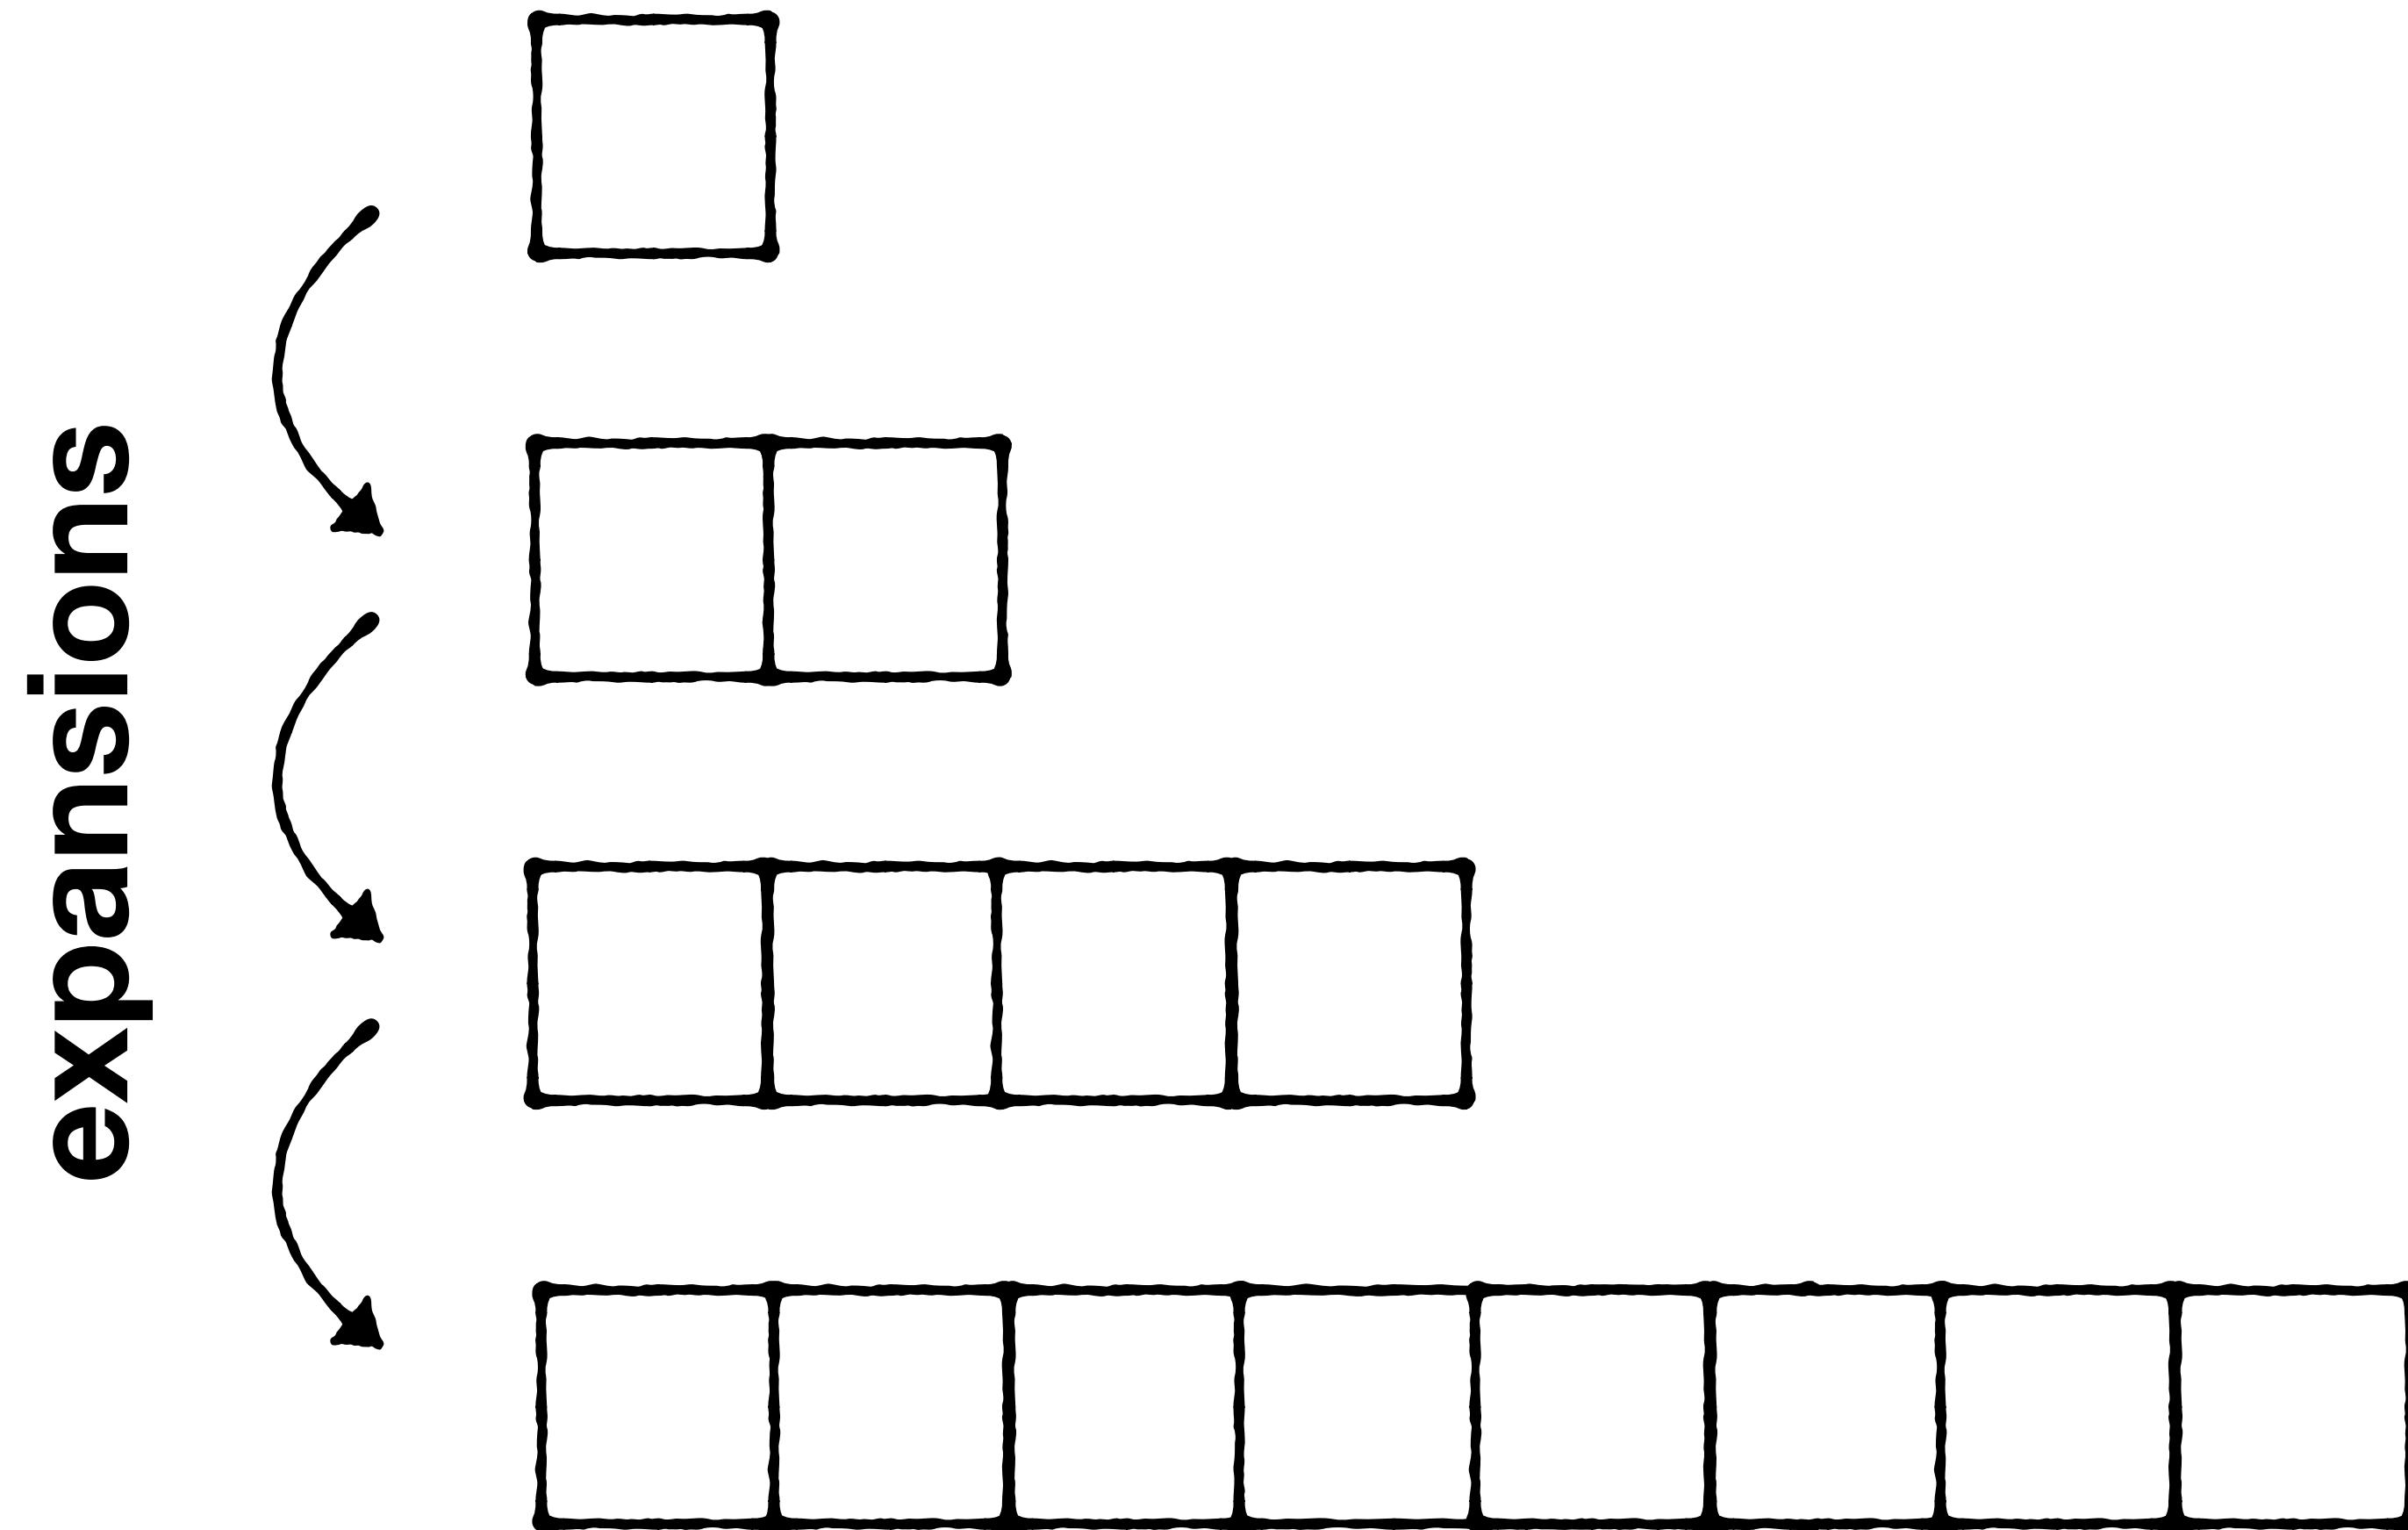


**Alleviating trade-
off via indirection**

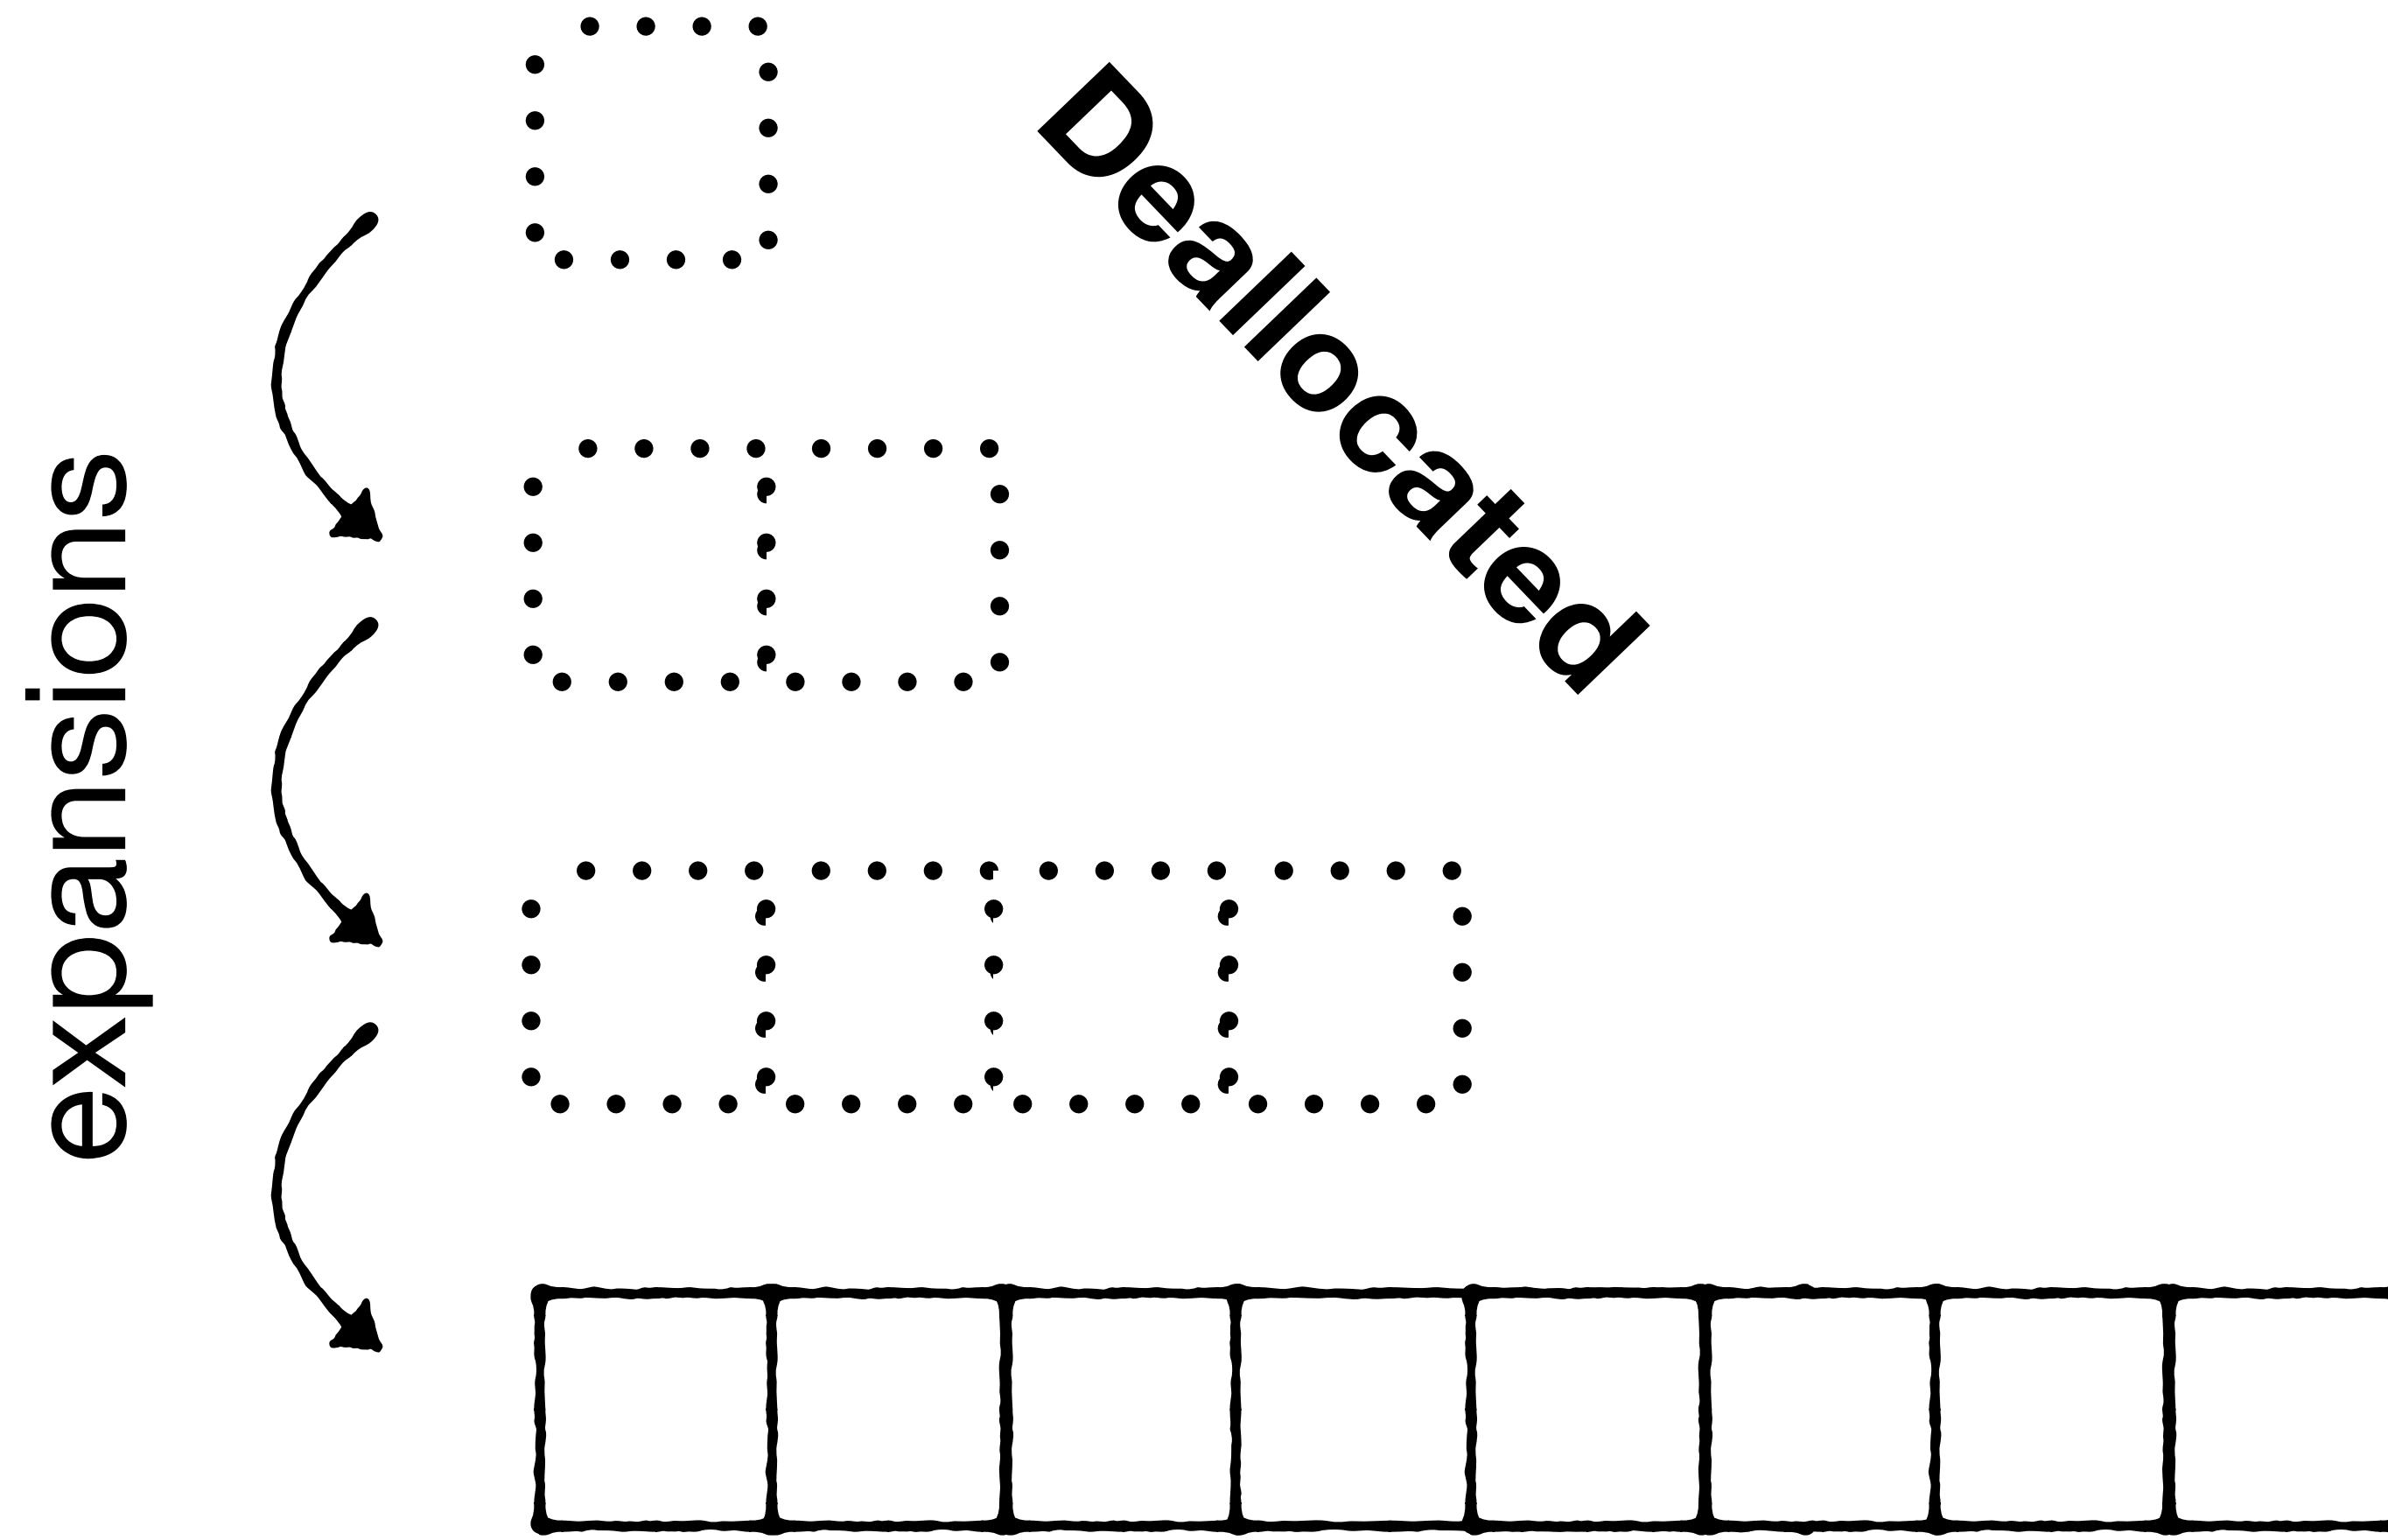


Reusing Deallocated Space

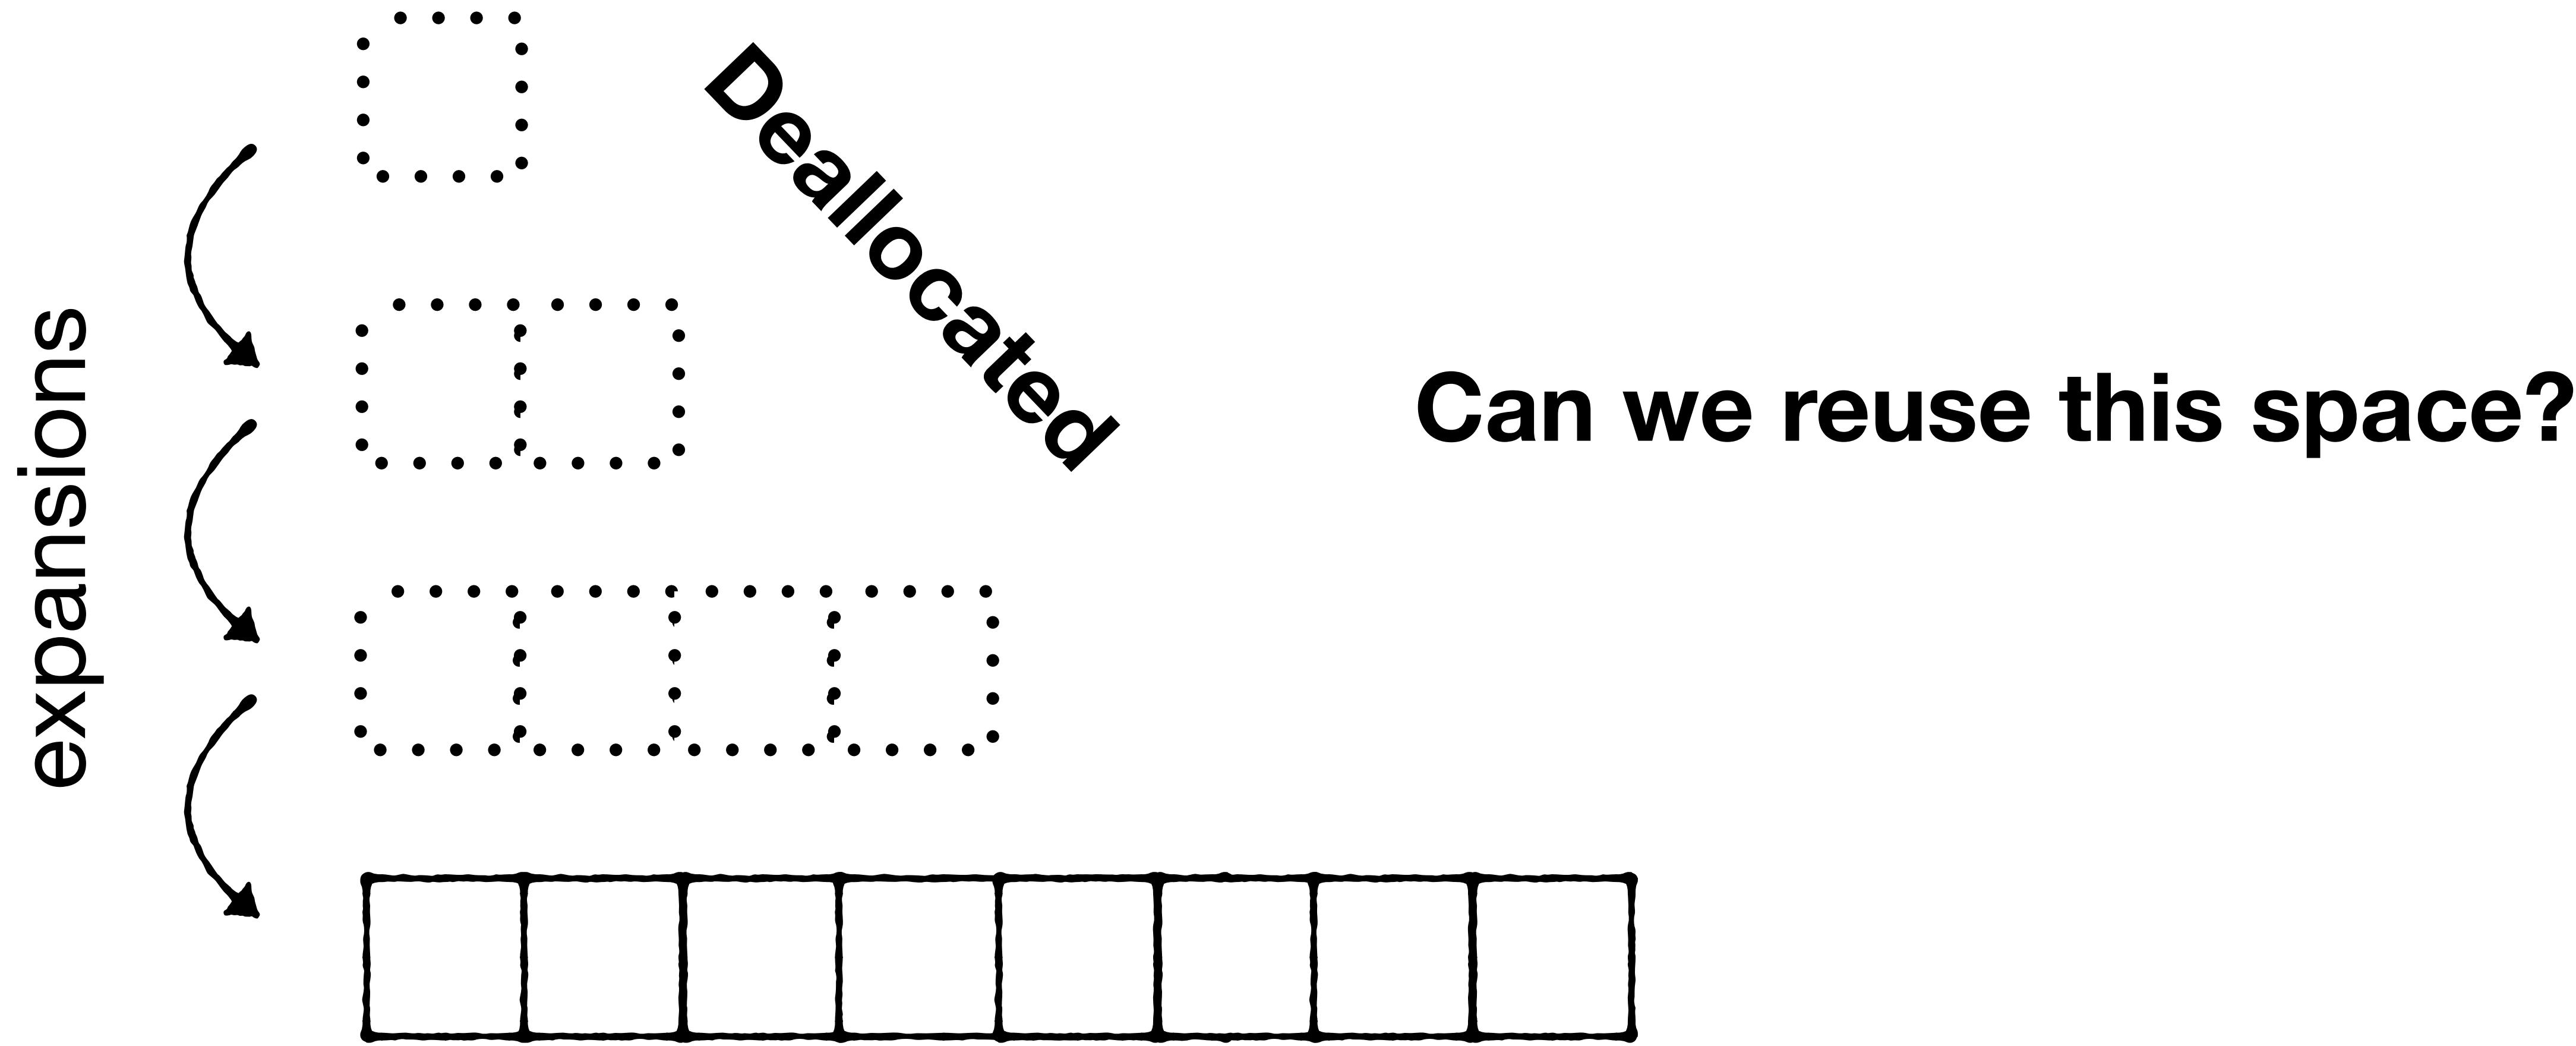
Assume $G \geq 2$



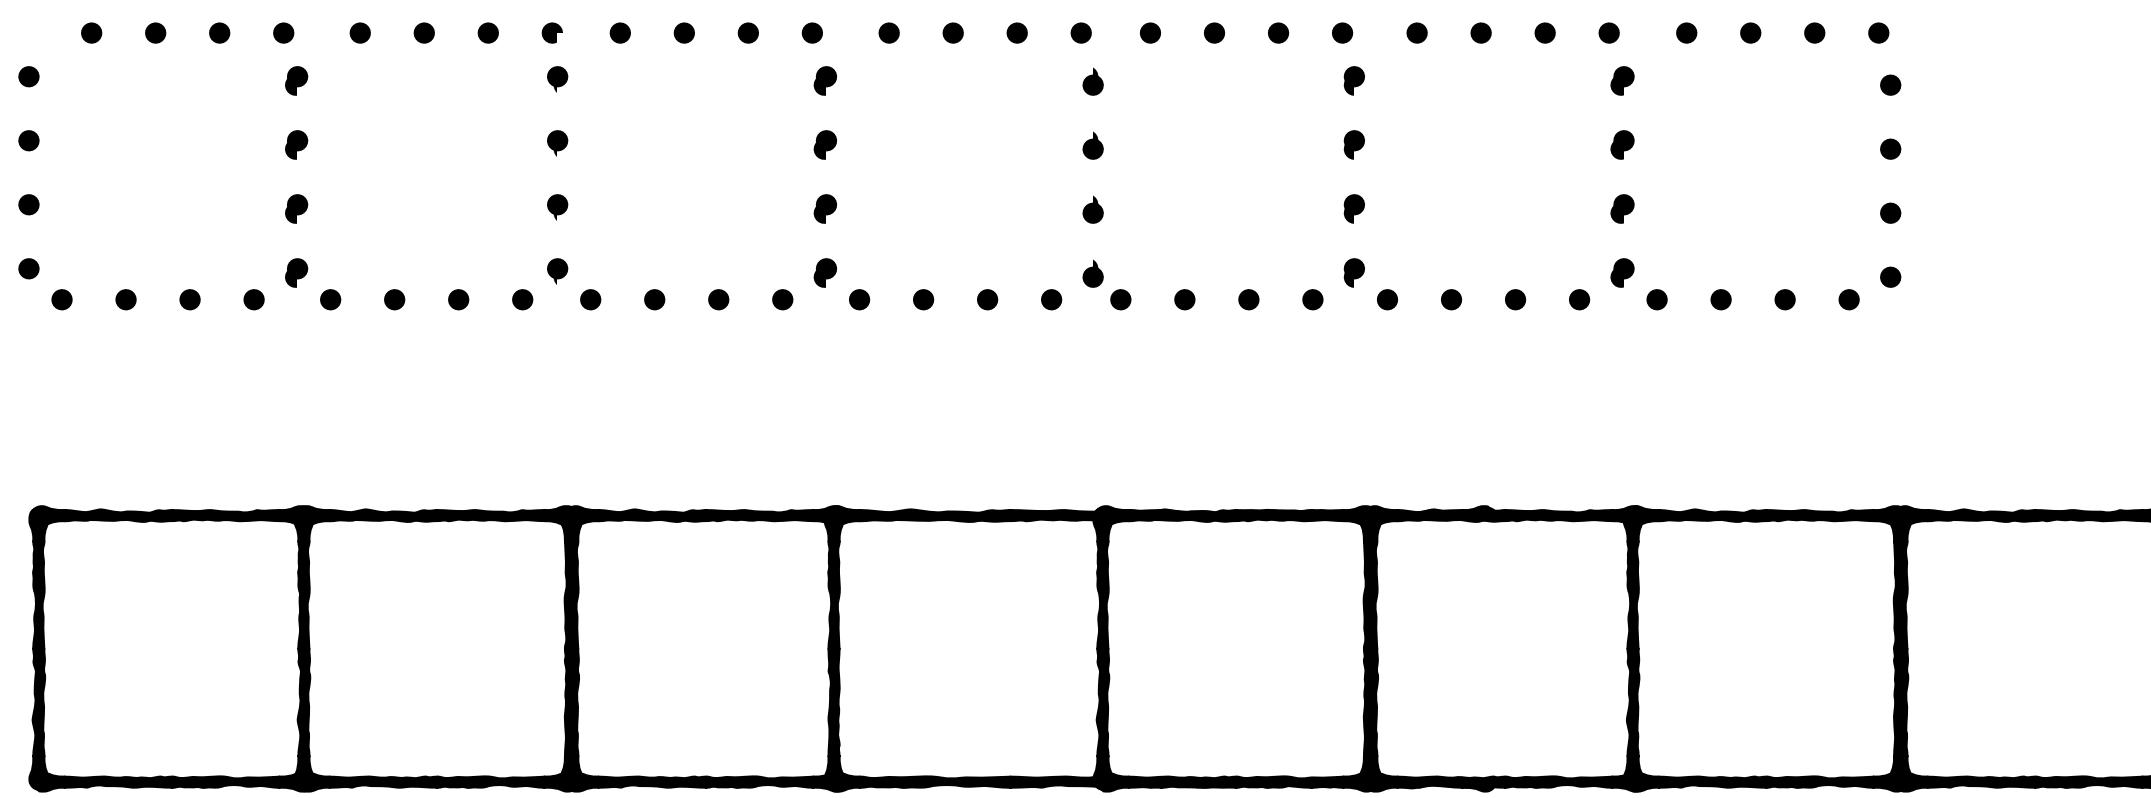
Reusing Deallocated Space ($G \geq 2$)



Reusing Deallocated Space ($G \geq 2$)



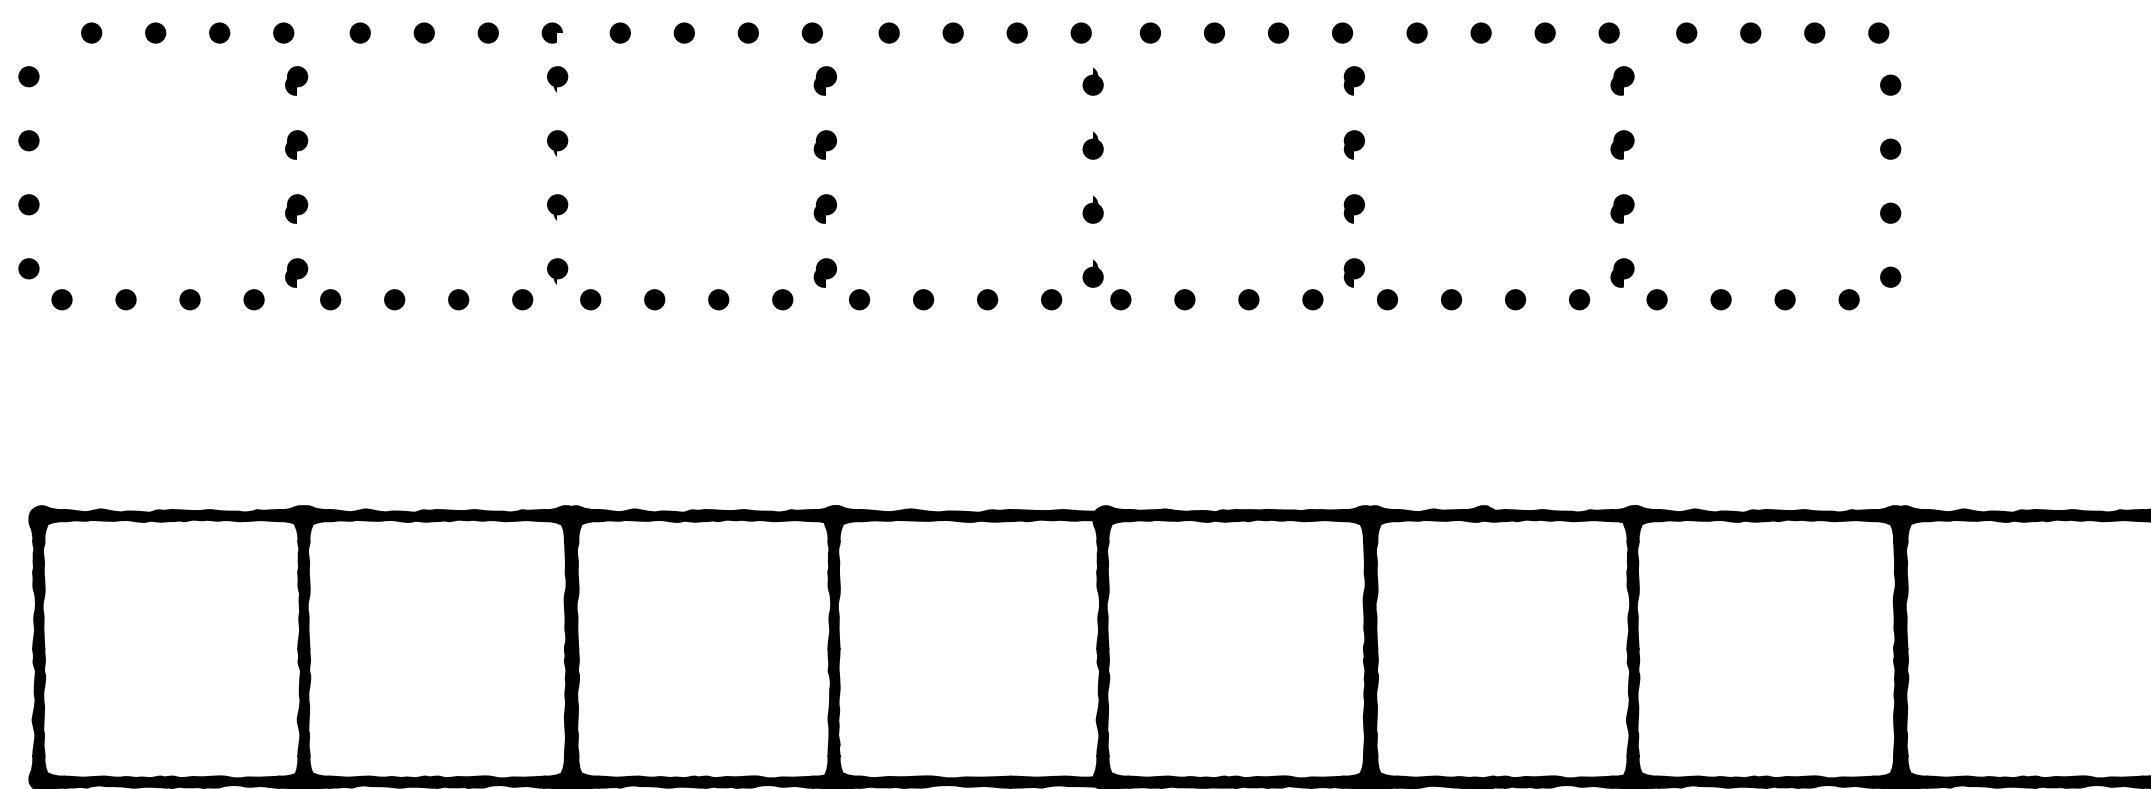
Reusing Deallocated Space ($G \geq 2$)



Total deallocated space
=

Total allocated space - 1

Reusing Deallocated Space ($G \geq 2$)



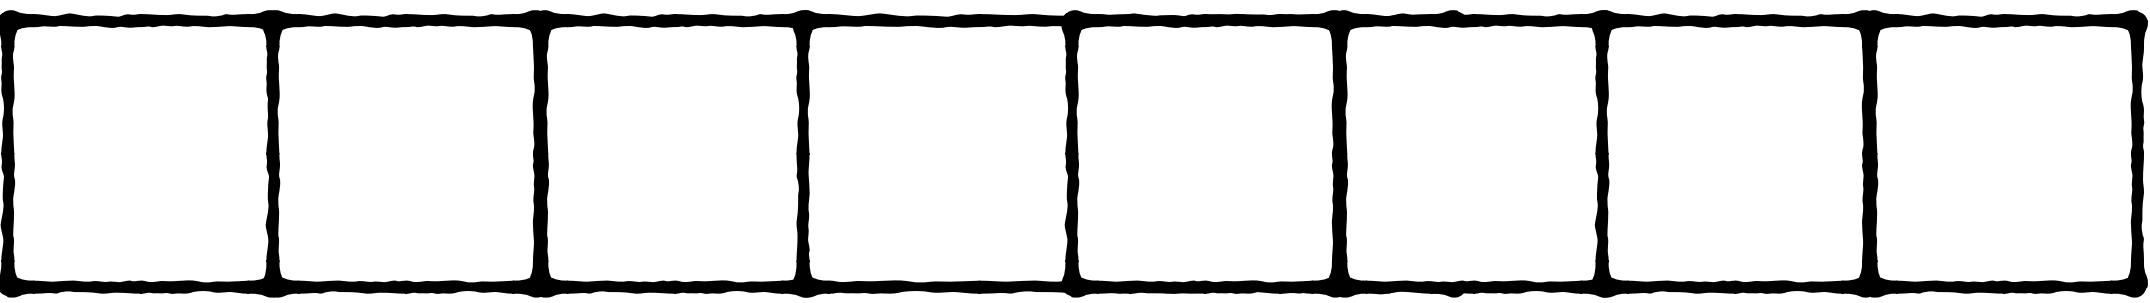
Total deallocated space

Λ

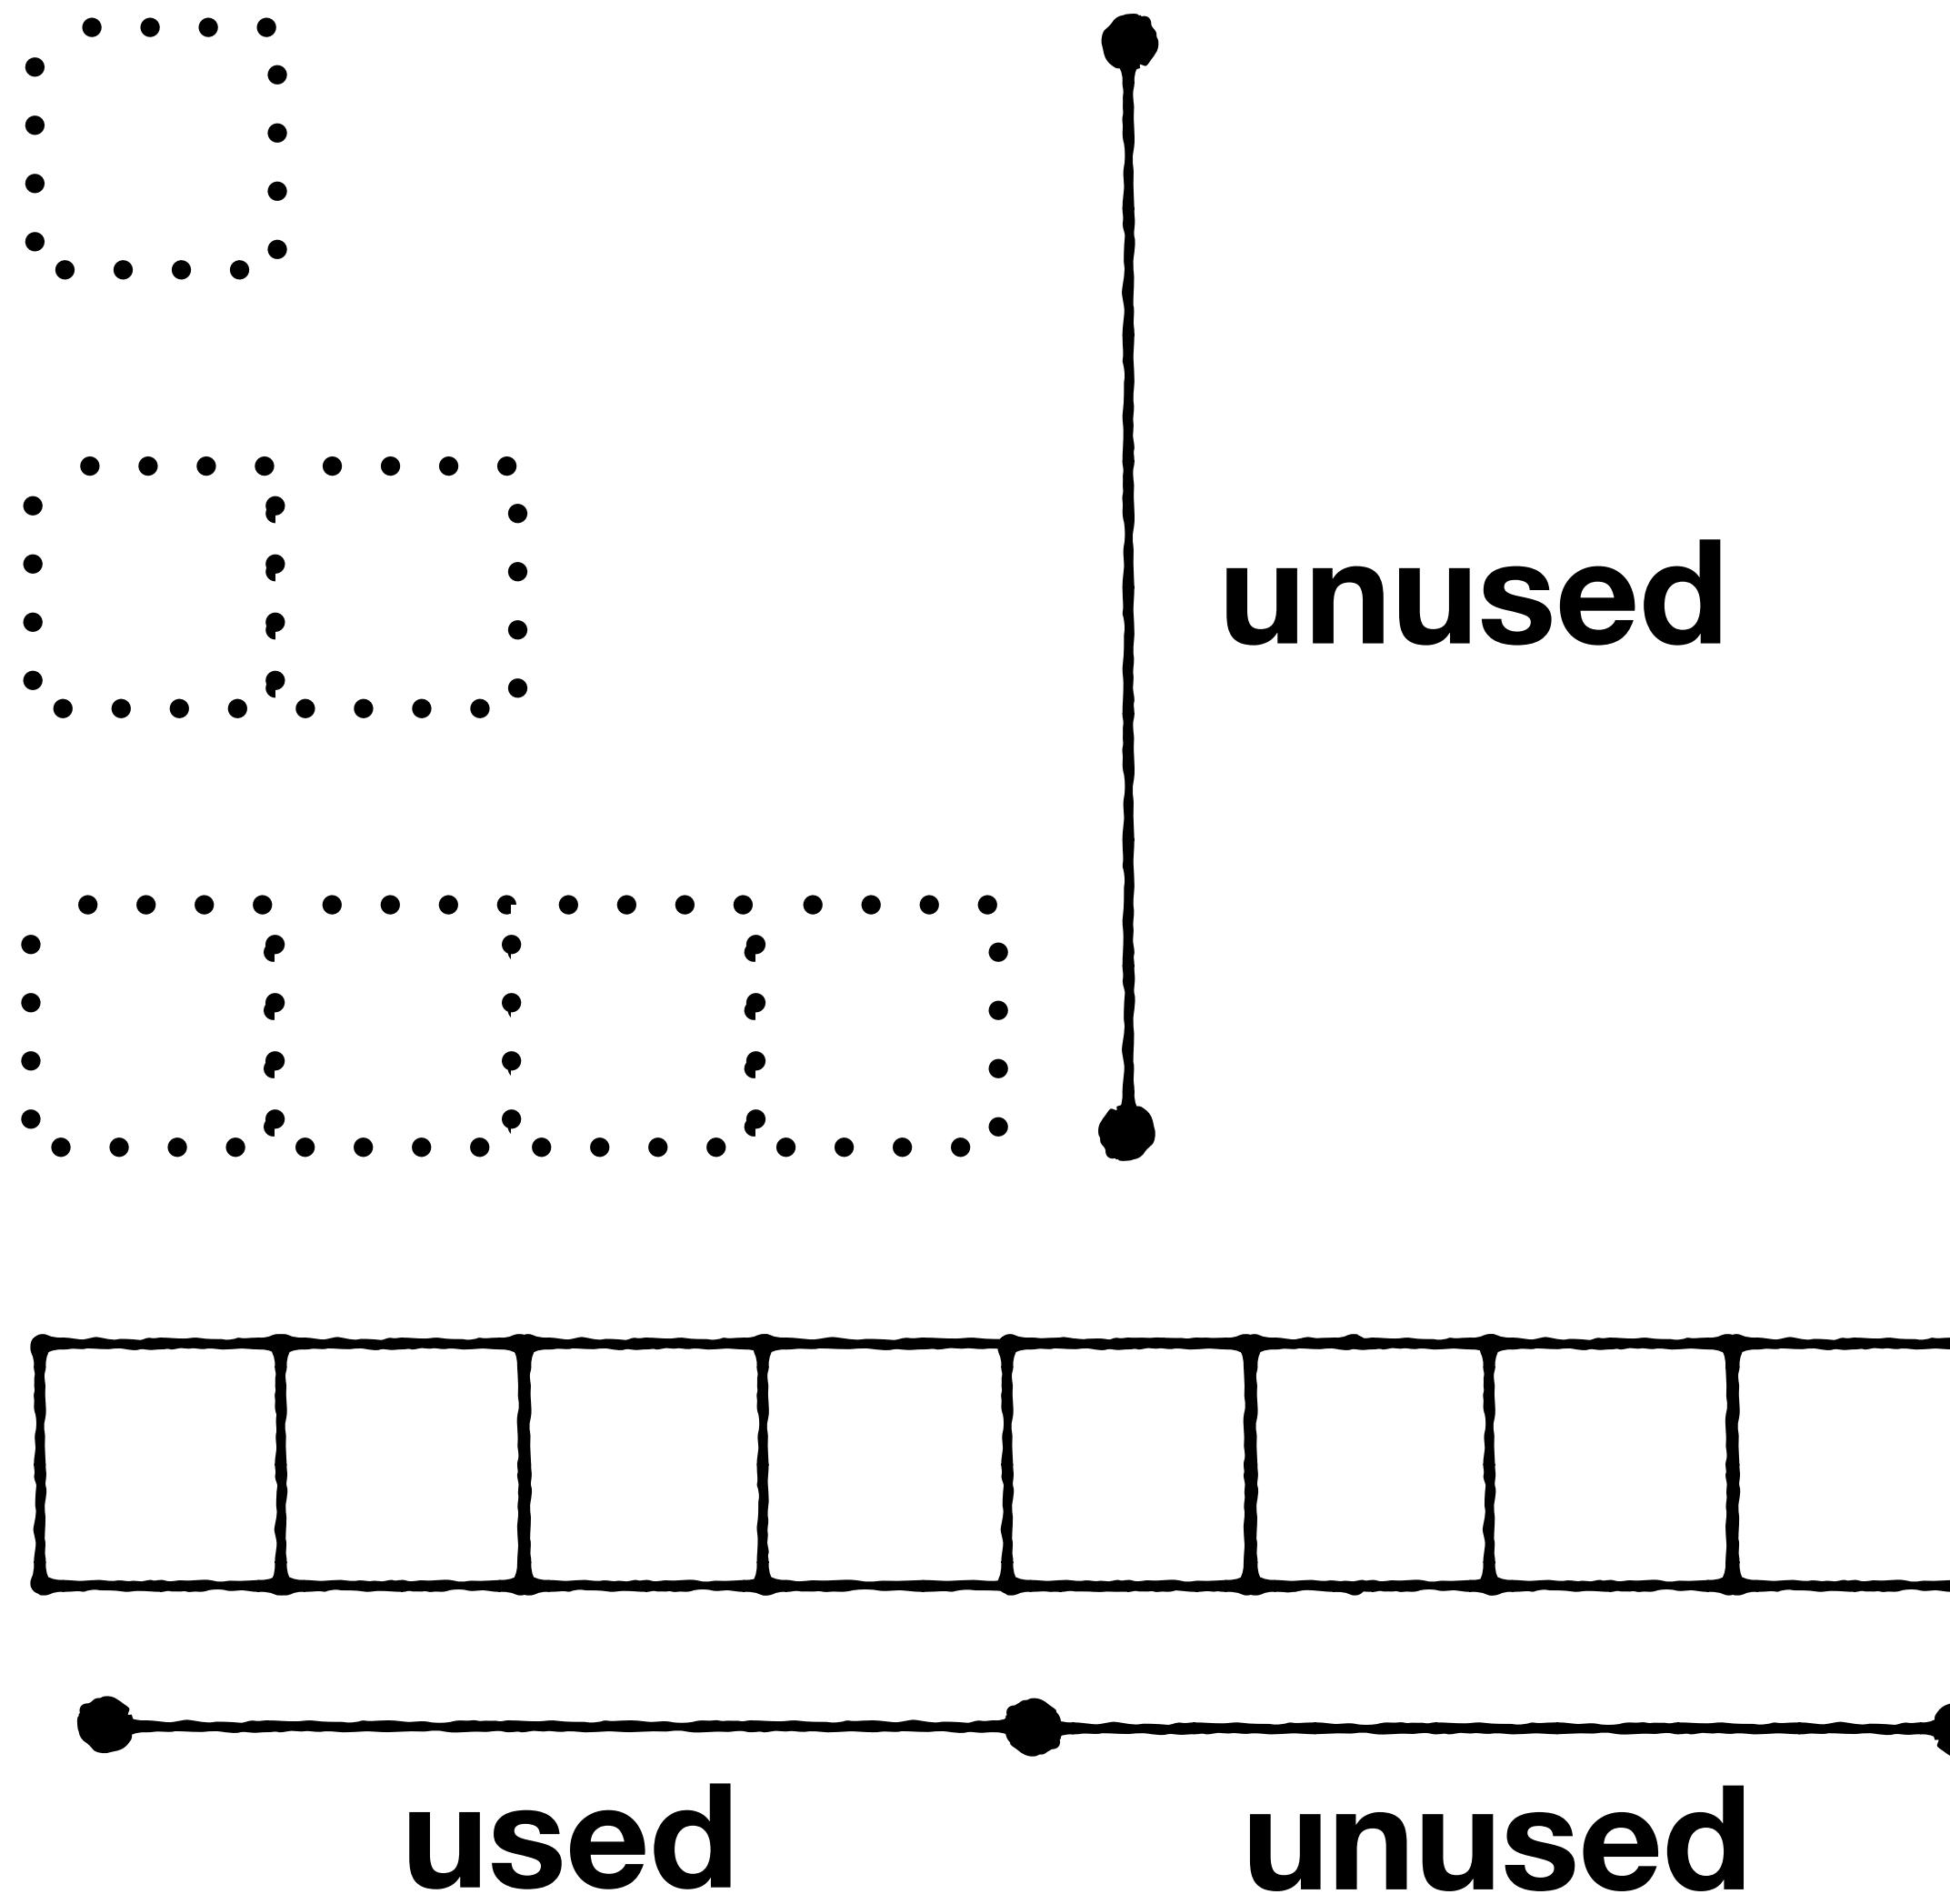
Total allocated space

Reusing Deallocated Space ($G \geq 2$)

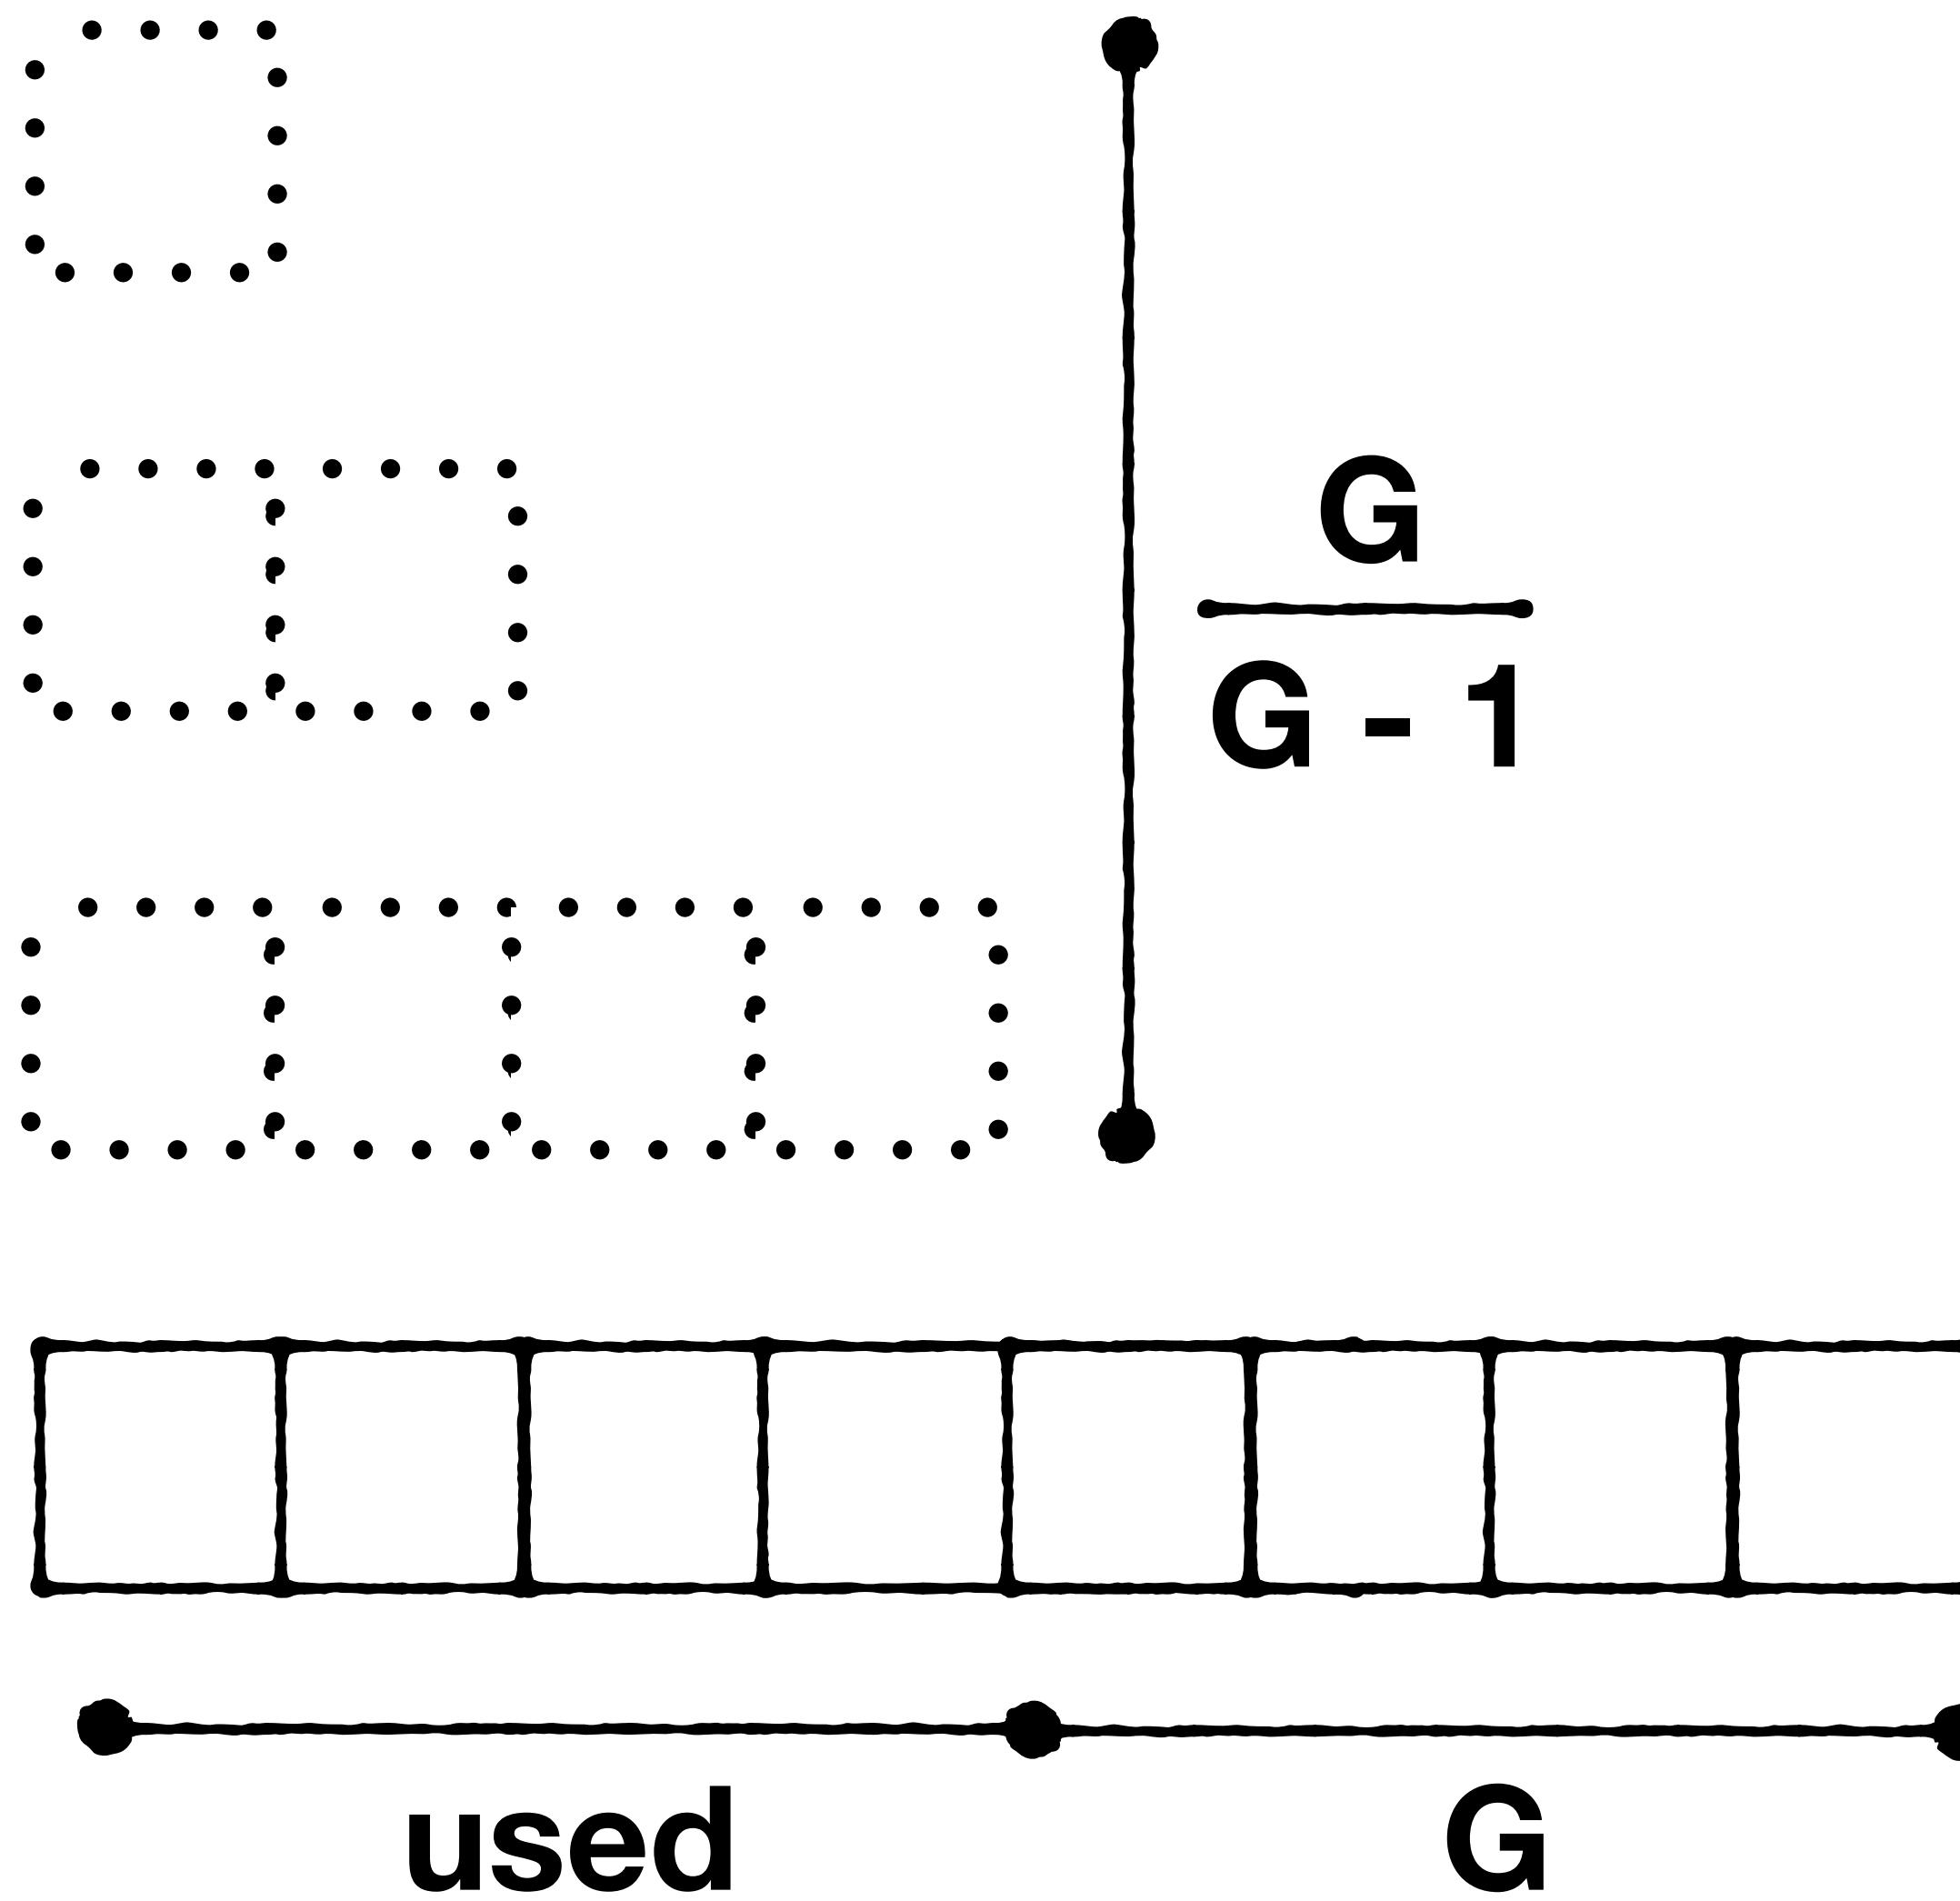
**Deallocated space
Can't be reused by new array**



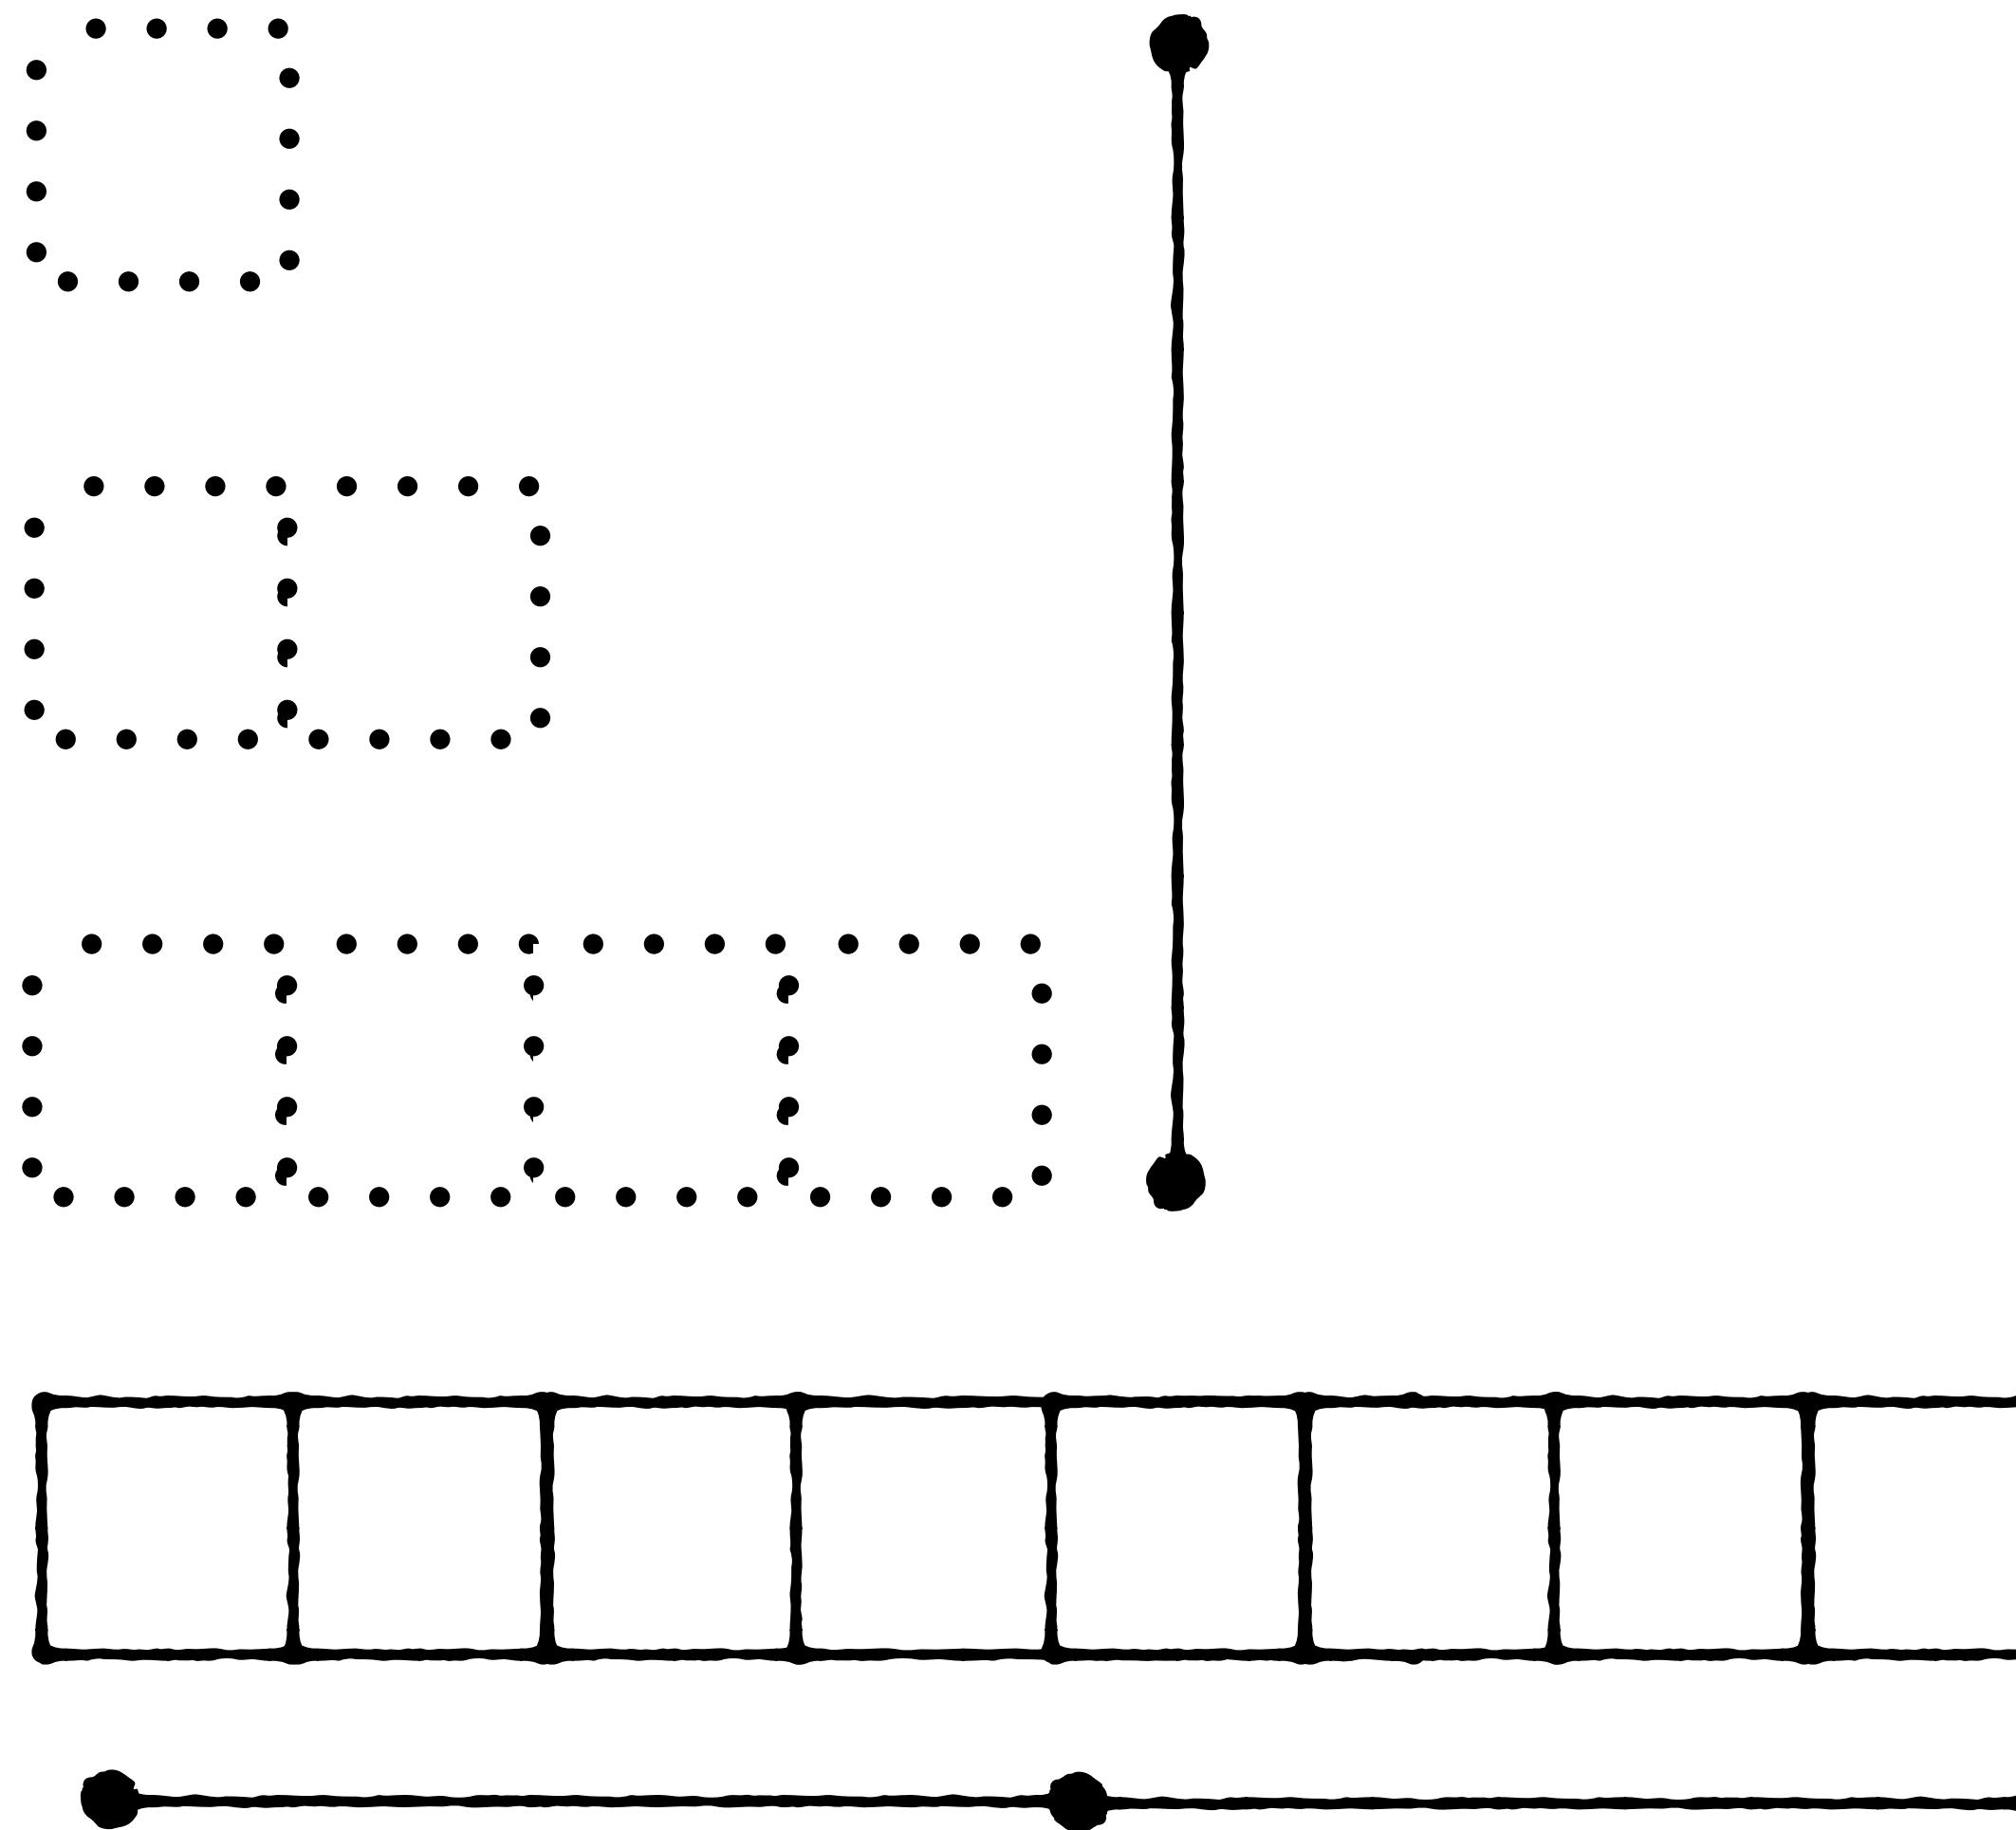
Deriving Max Space-Amp



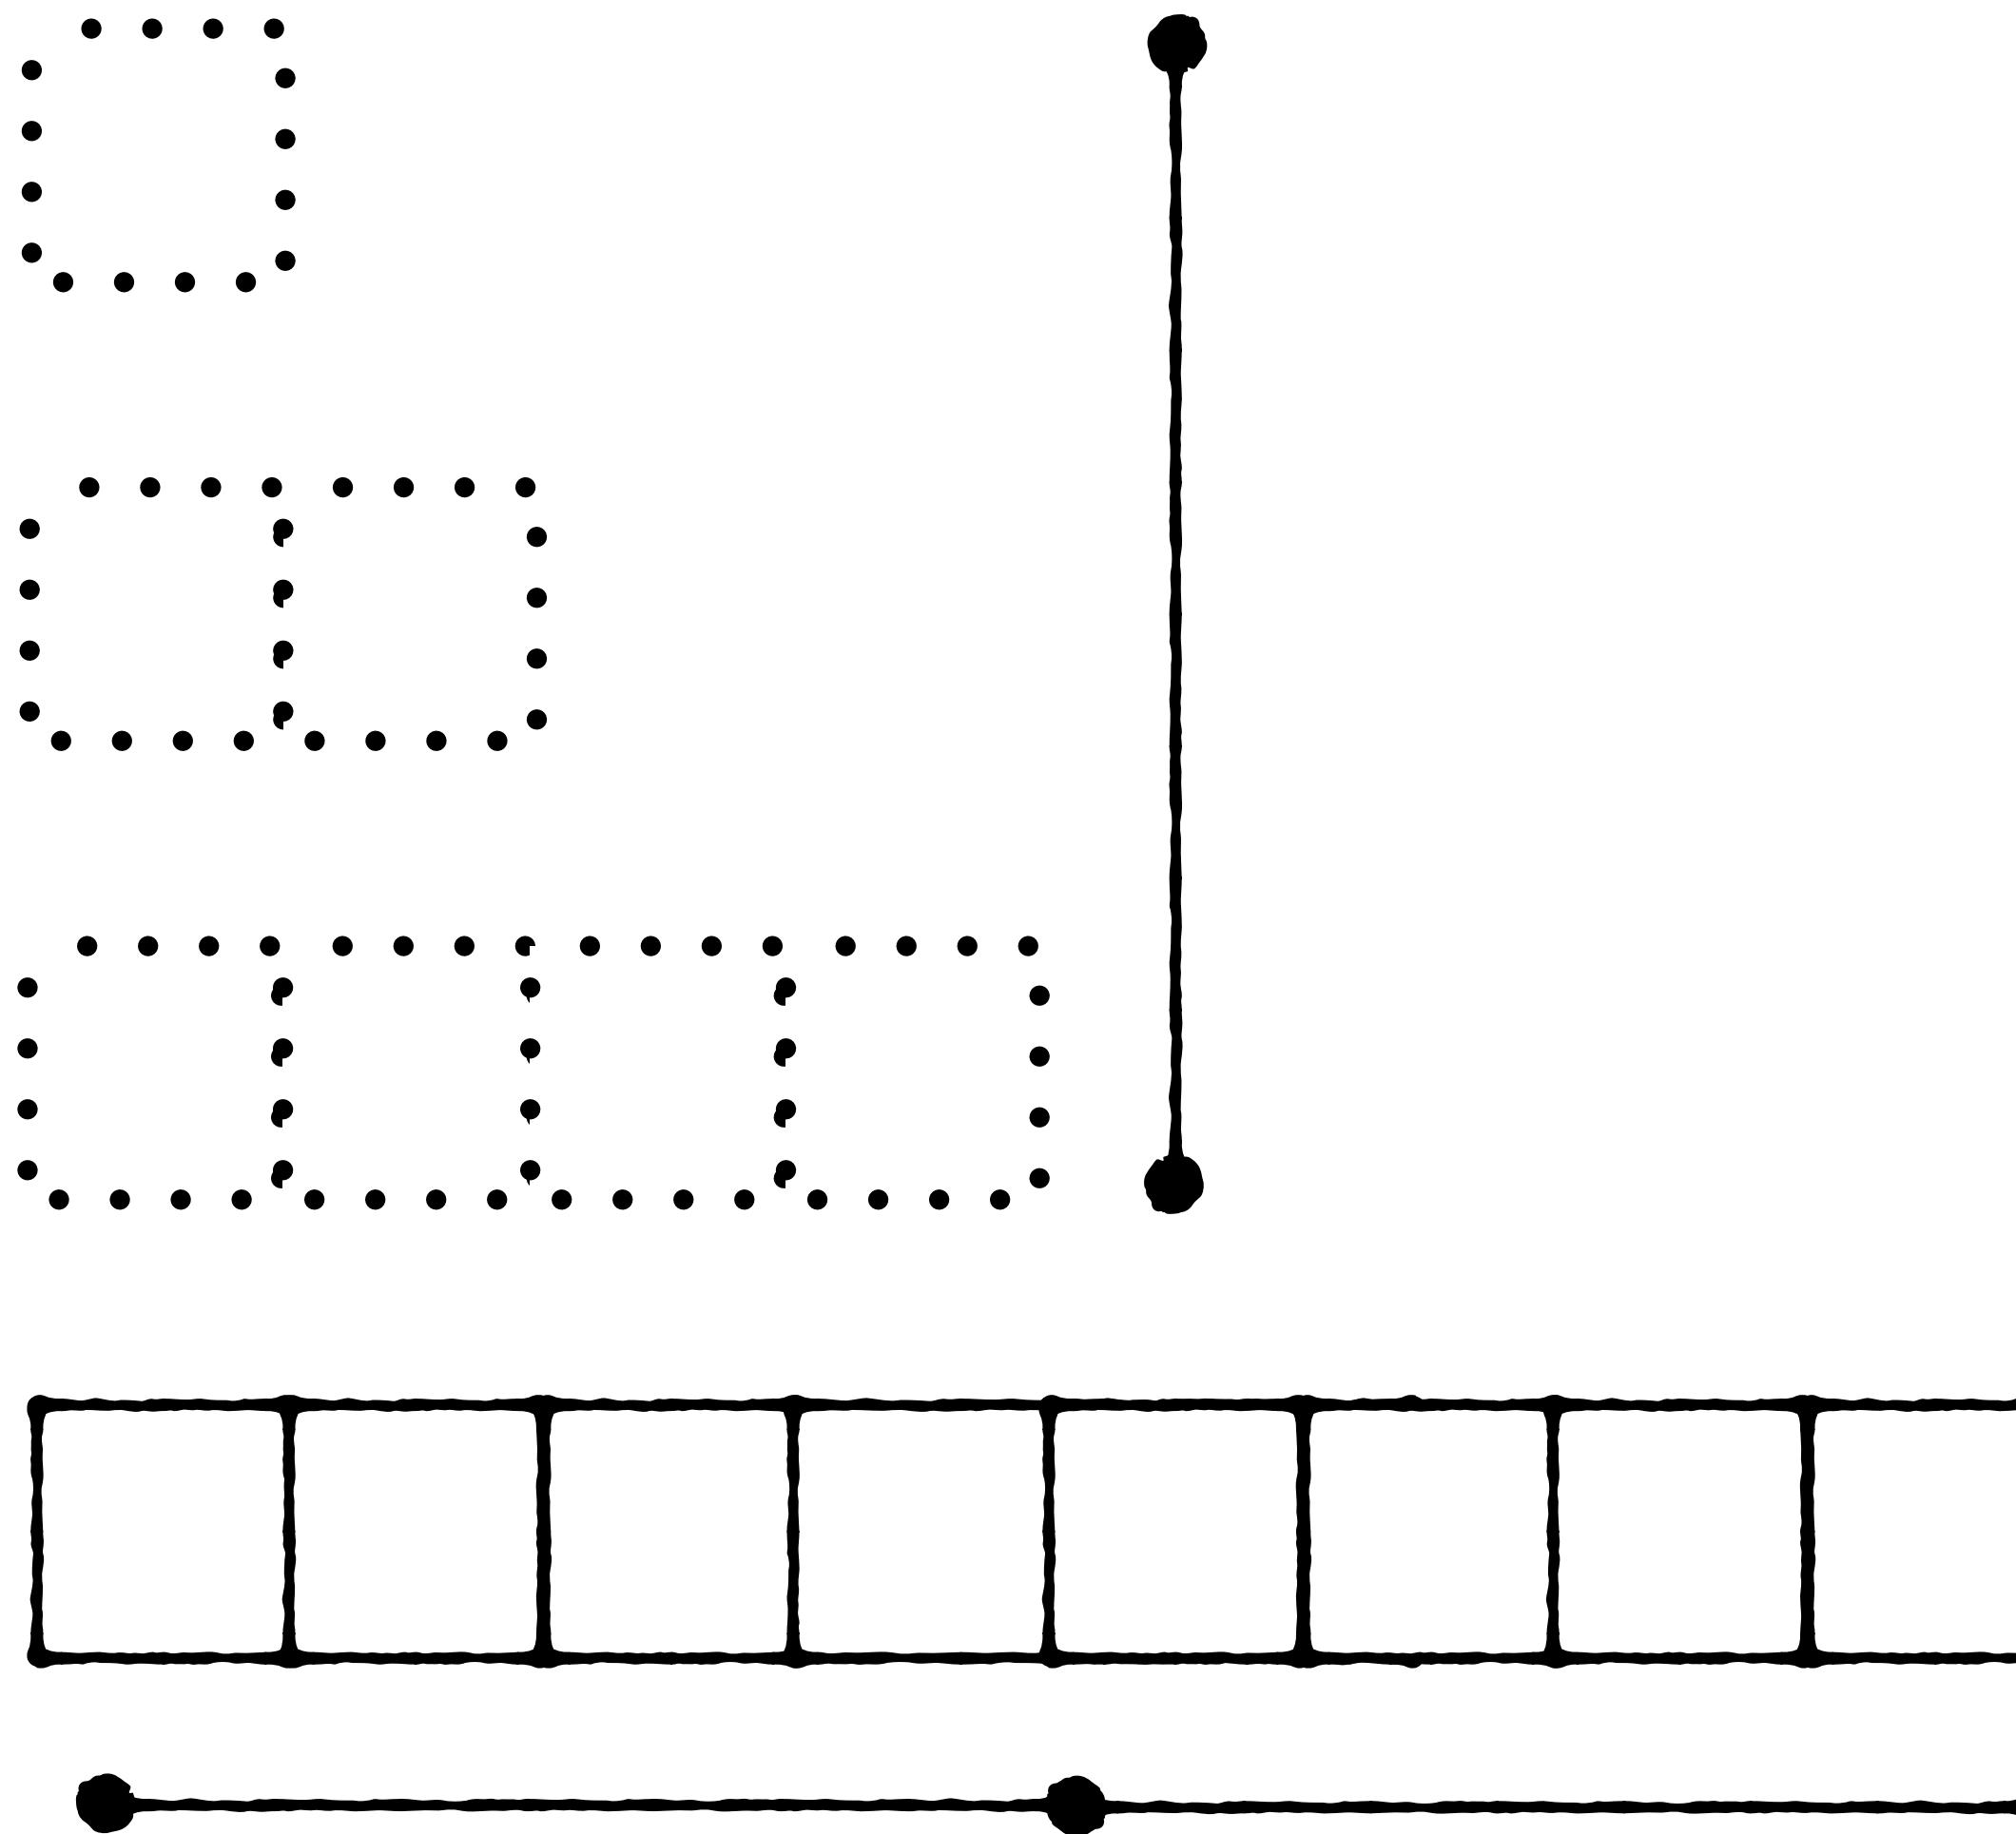
Deriving Max Space-Amp



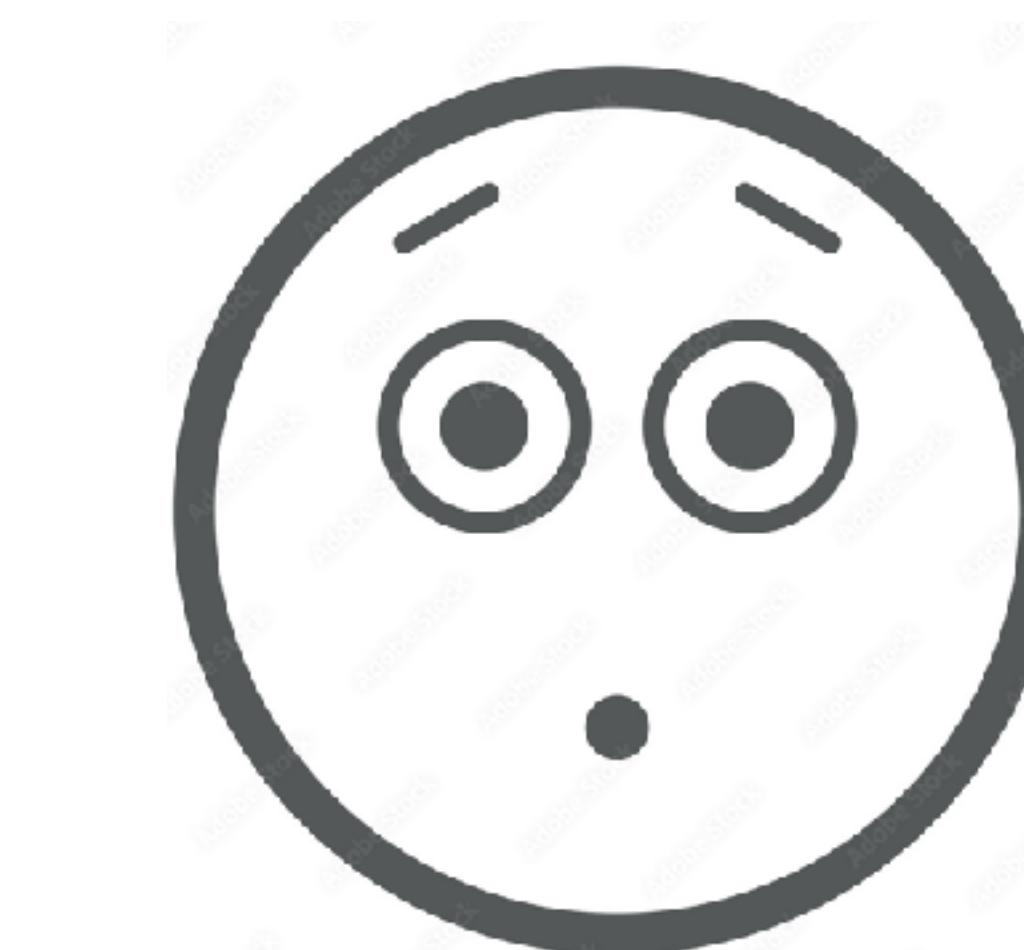
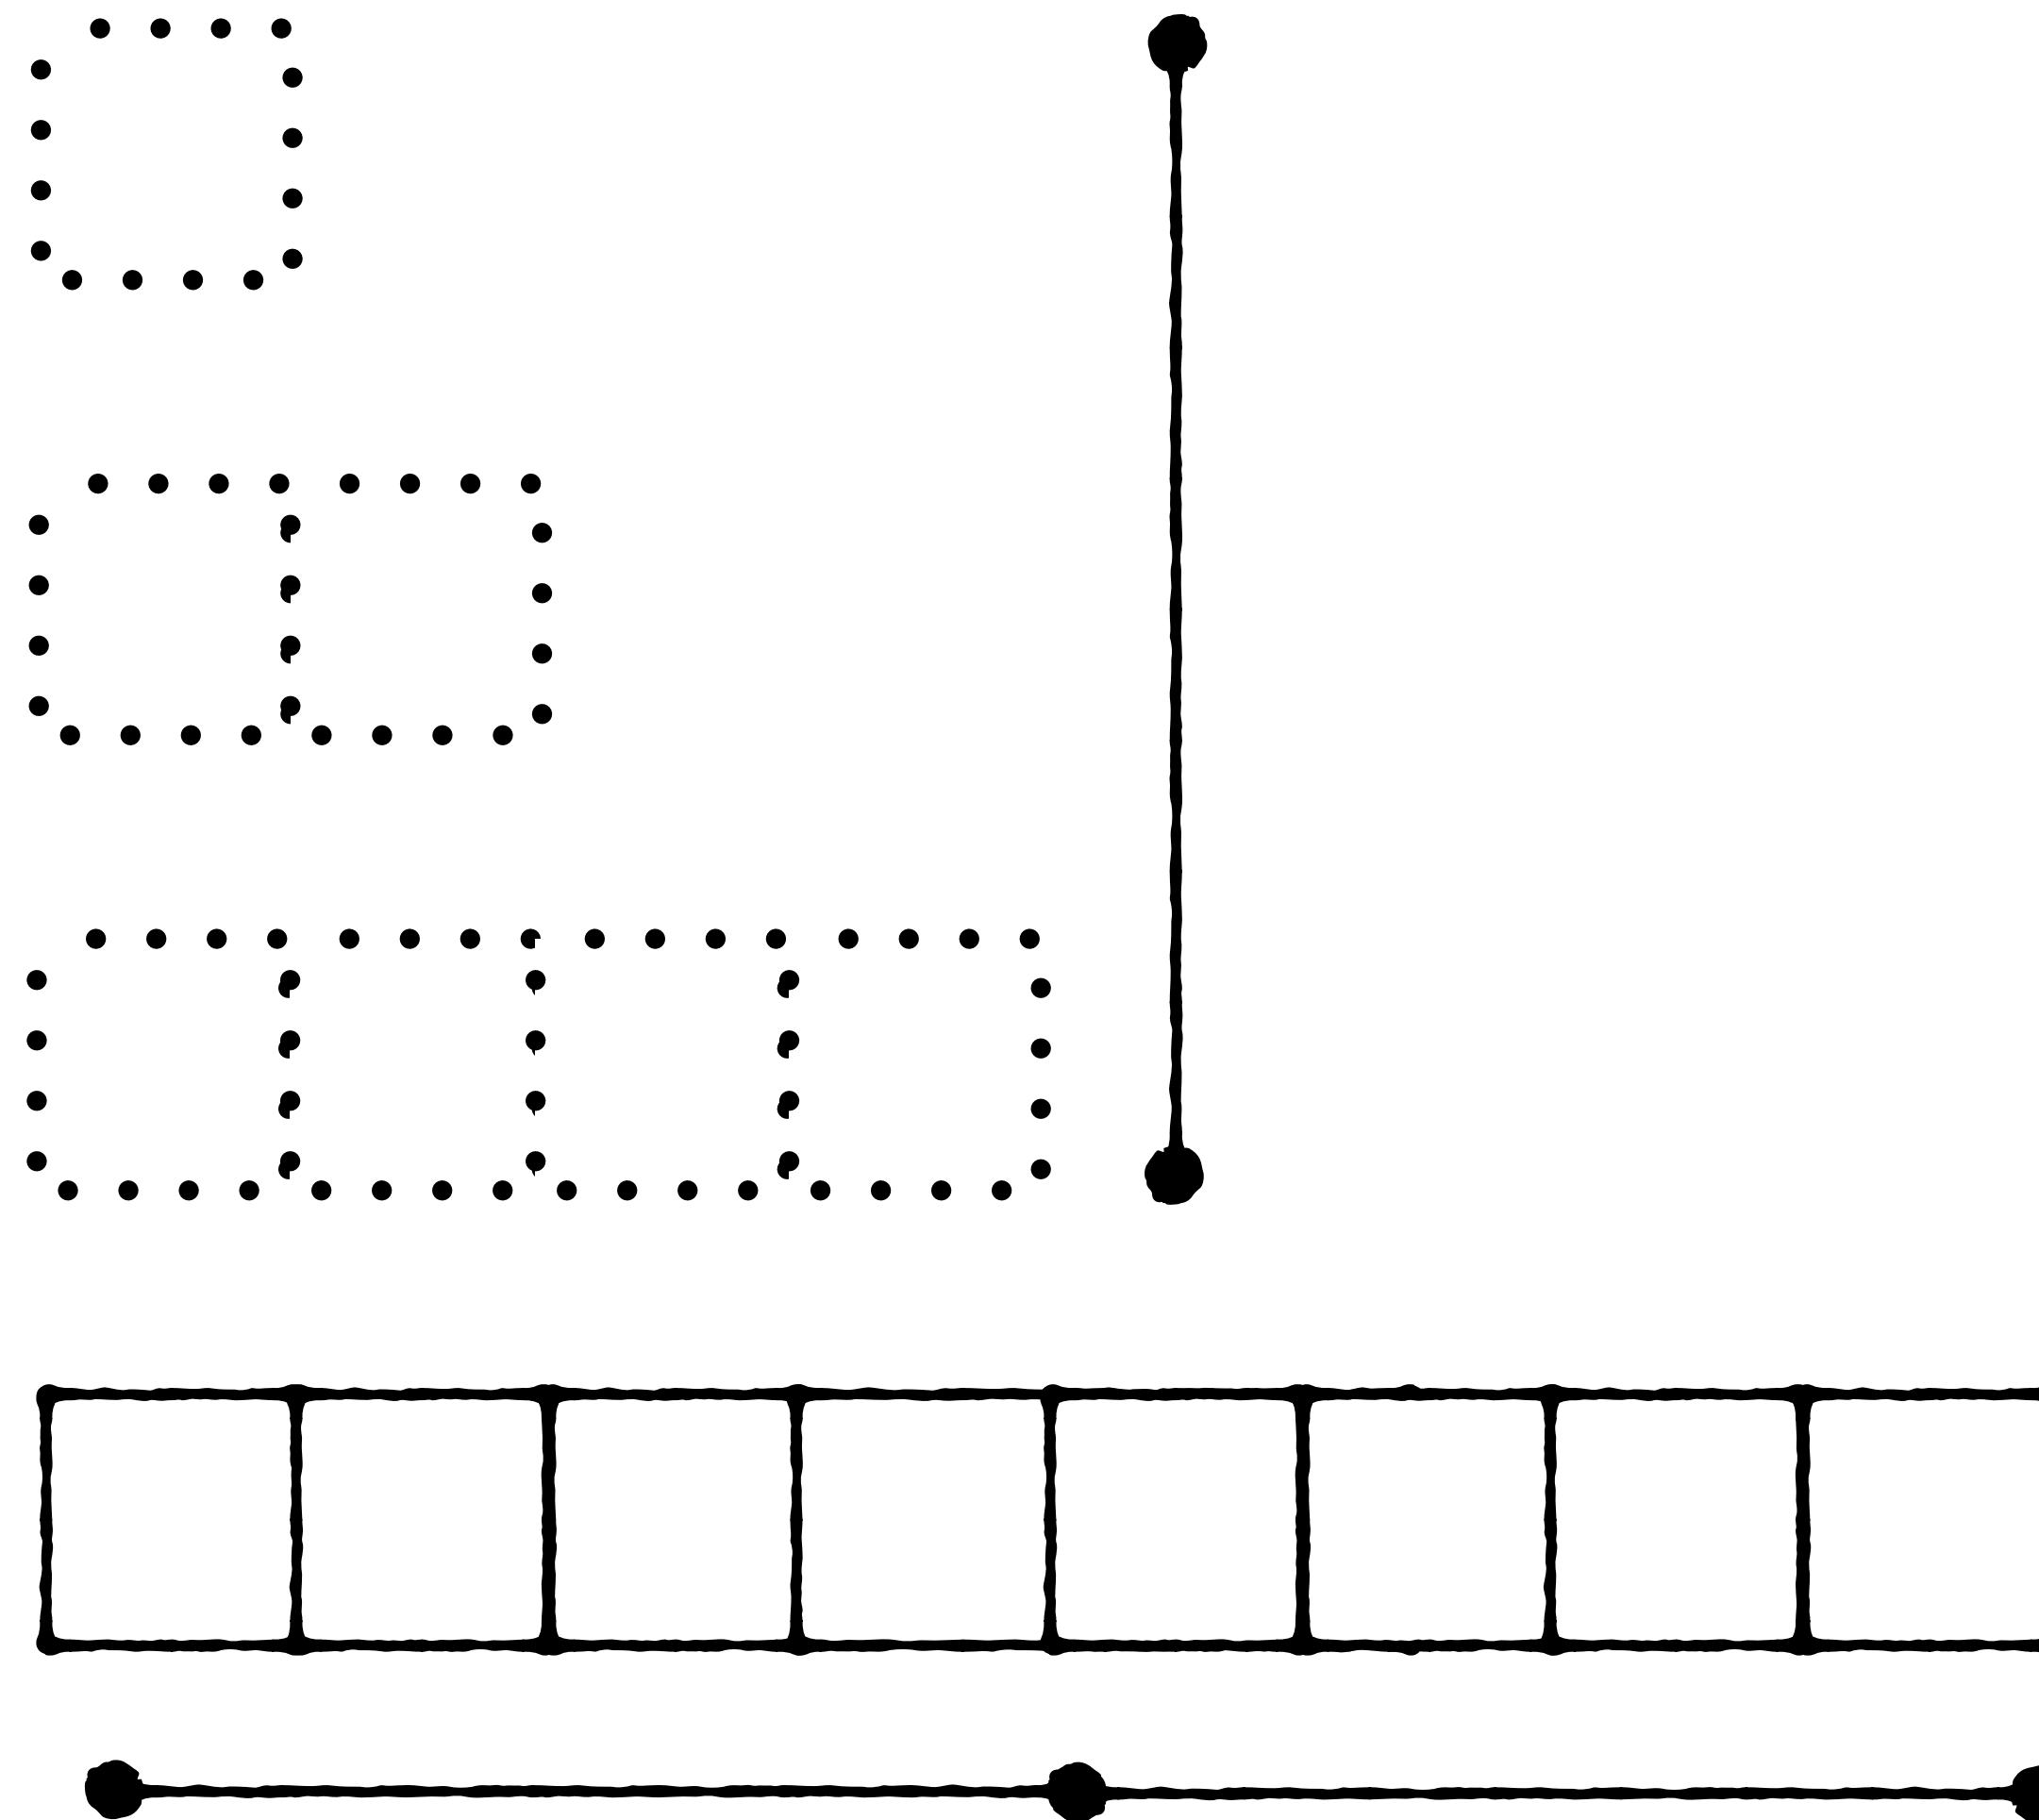
$$\text{Max Space-Amp} = \frac{\text{used} + \text{unused}}{\text{used}} = \frac{G}{G - 1} + G$$



$$\text{Max Space-Amp} = \frac{\text{used} + \text{unused}}{\text{used}} = \frac{G^2}{G - 1}$$

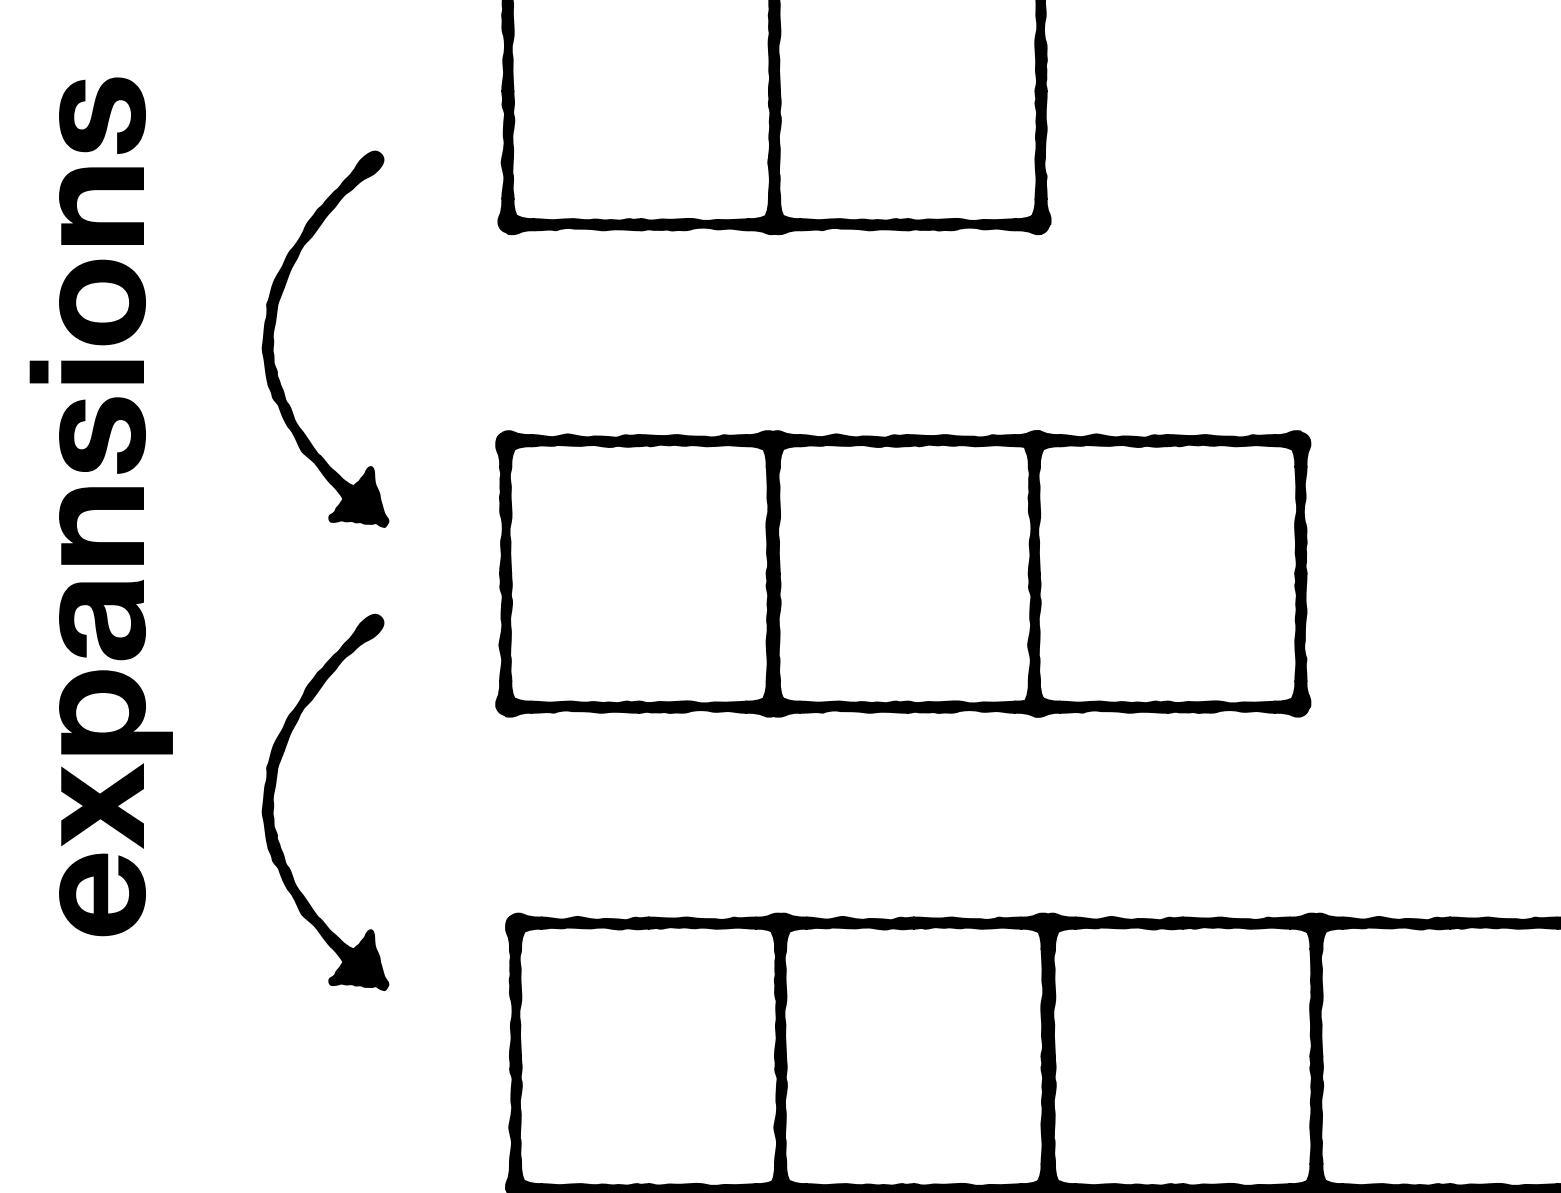


$$\text{Max Space-Amp} = \frac{\text{used} + \text{unused}}{\text{used}} = \frac{G^2}{G - 1} = 4 \quad \text{for } G = 2$$

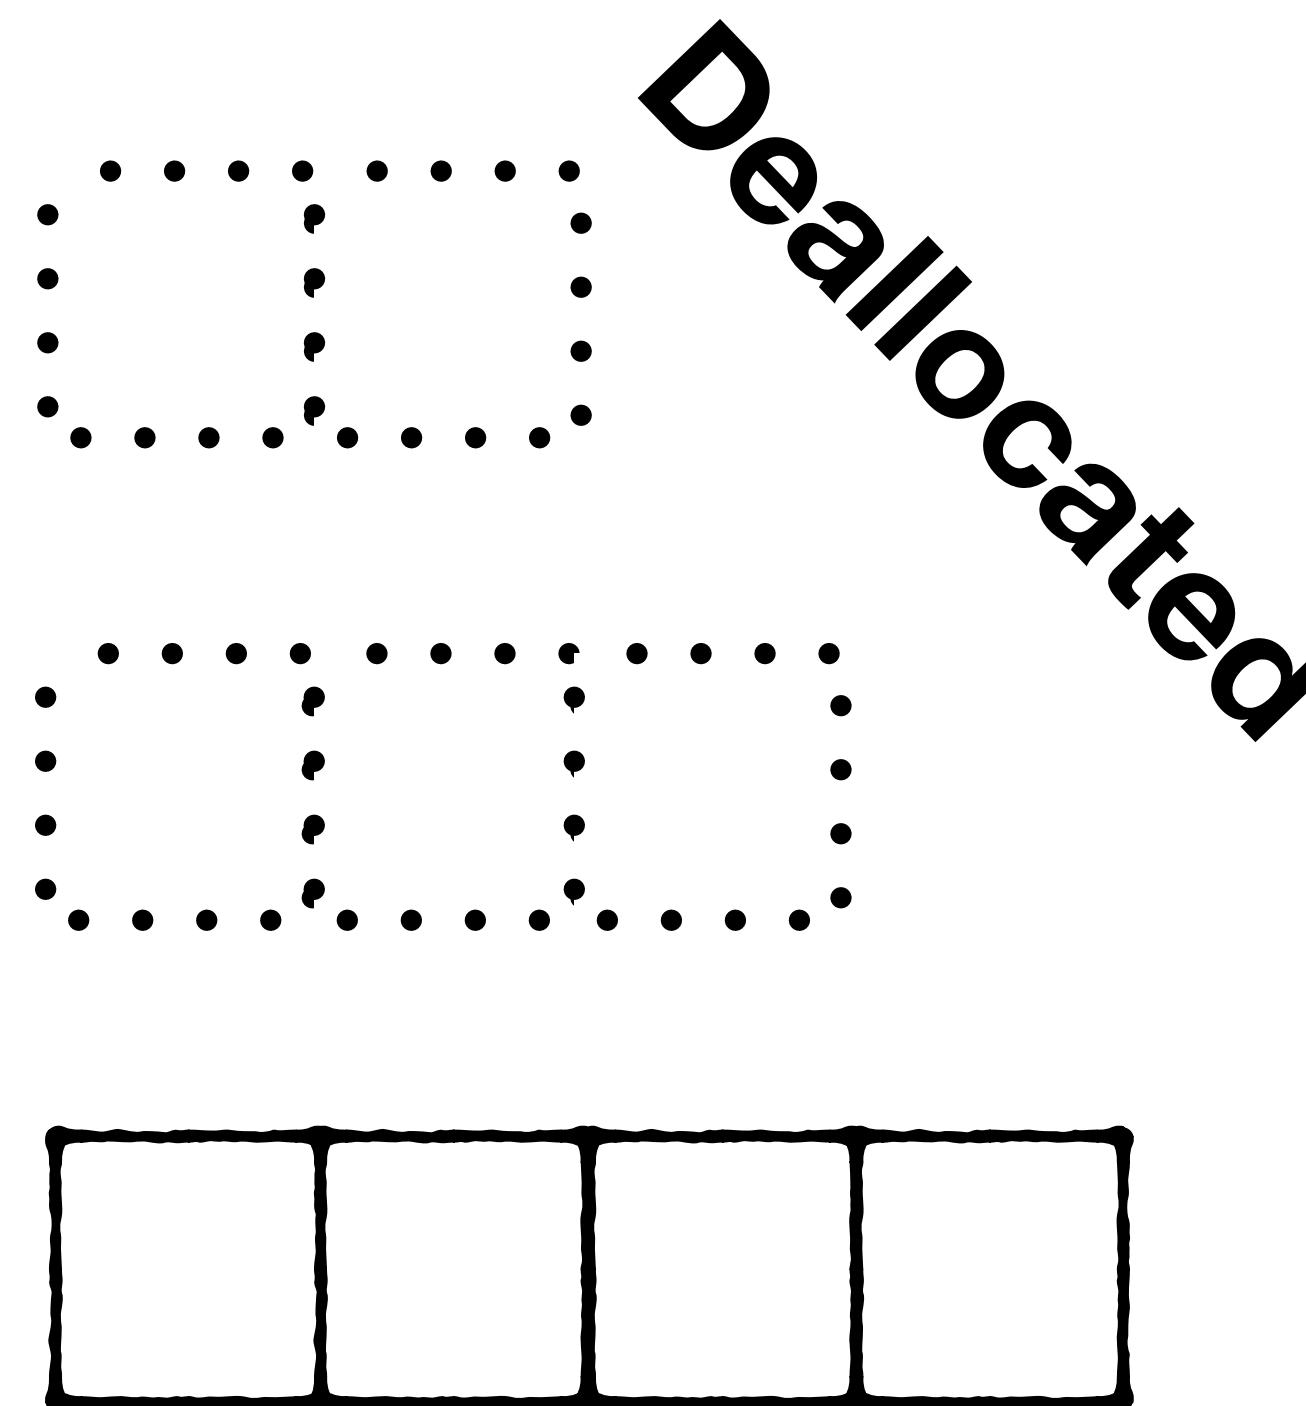


**Suppose we use small size ratio,
e.g., $G = 1.2$**

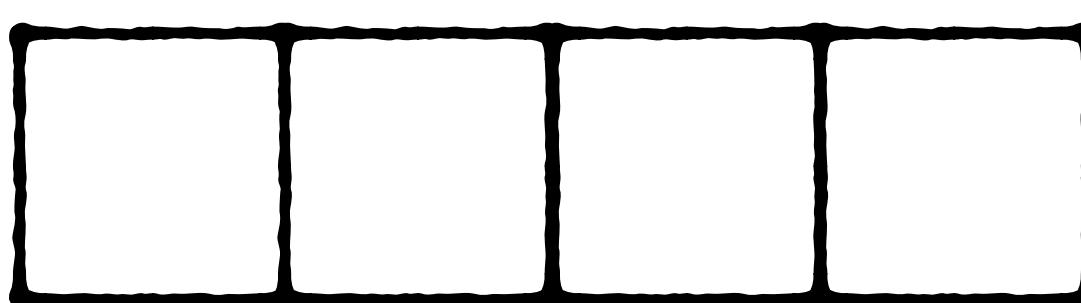
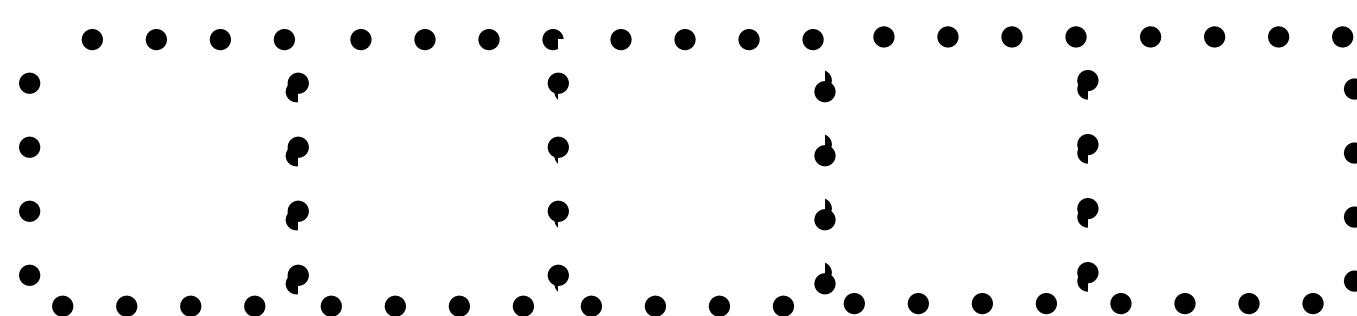
Suppose we use small size ratio,
e.g., $G = 1.2$



Suppose we use small size ratio,
e.g., $G = 1.2$



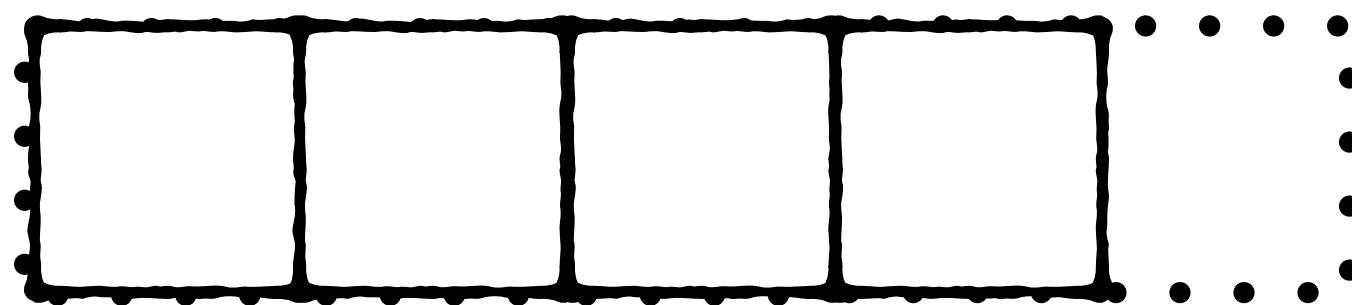
Suppose we use small size ratio,
e.g., $G = 1.2$



Total deallocated space

Total allocated space

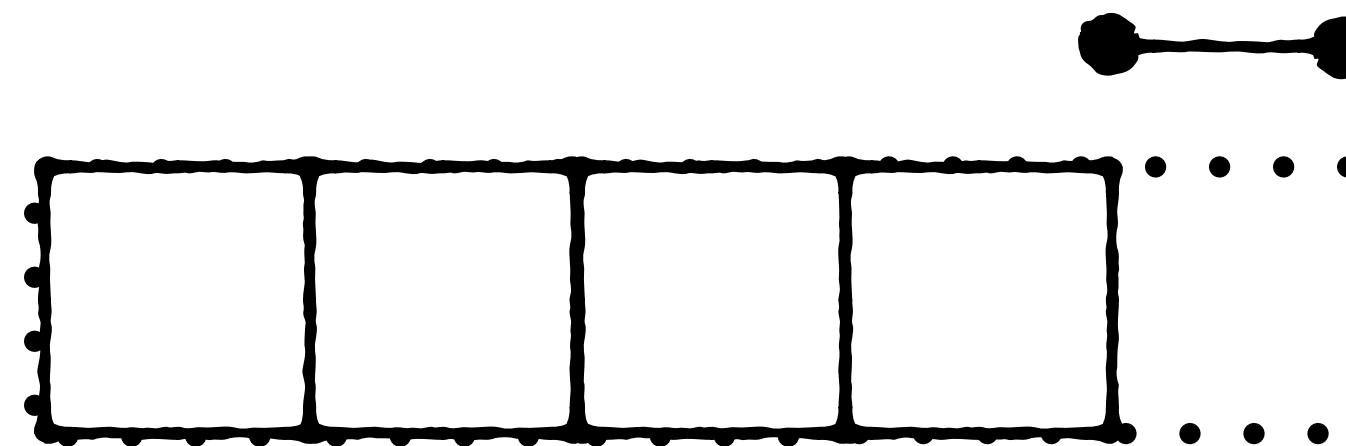
**Suppose we use small size ratio,
e.g., $G = 1.2$**



Reuse is possible

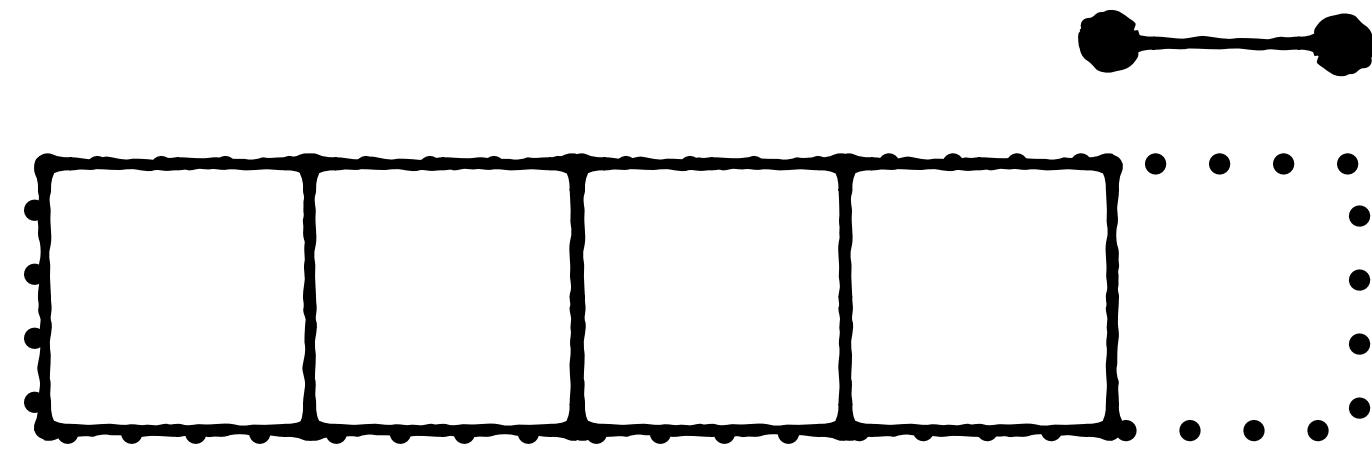
Suppose we use small size ratio,
e.g., $G = 1.2$

Wasted space



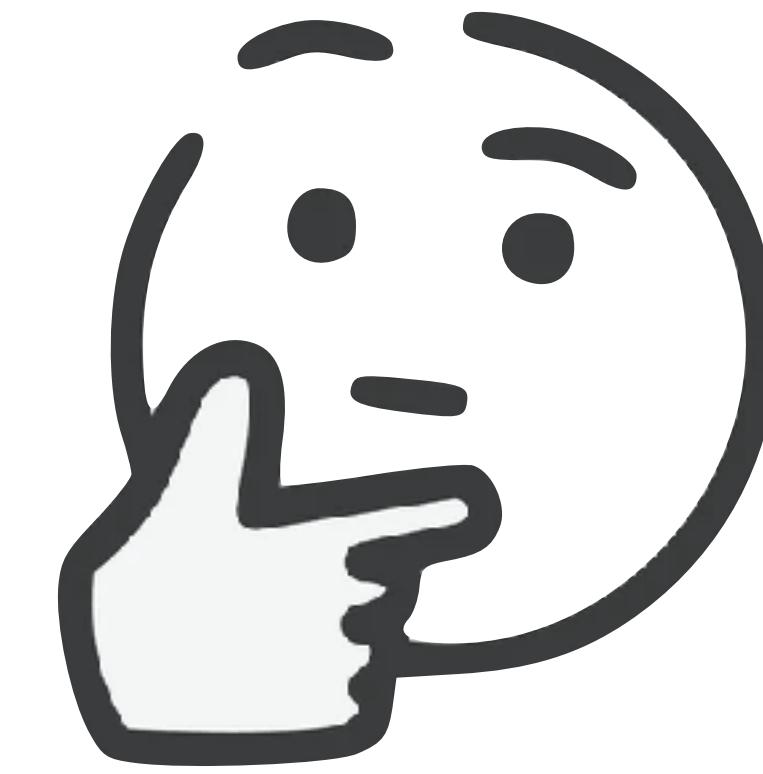
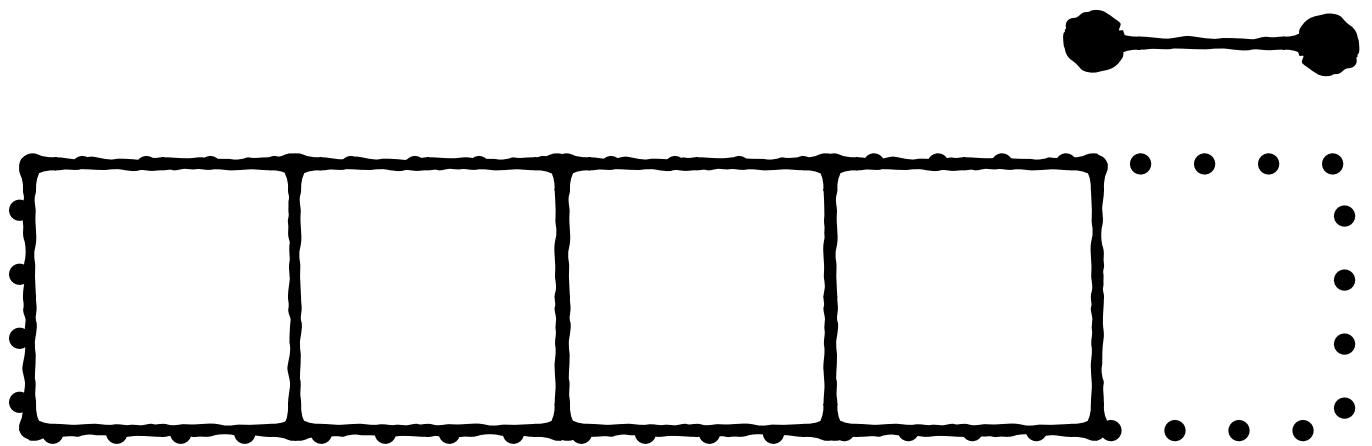
Suppose we use small size ratio,
e.g., $G = 1.2$

Wasted space

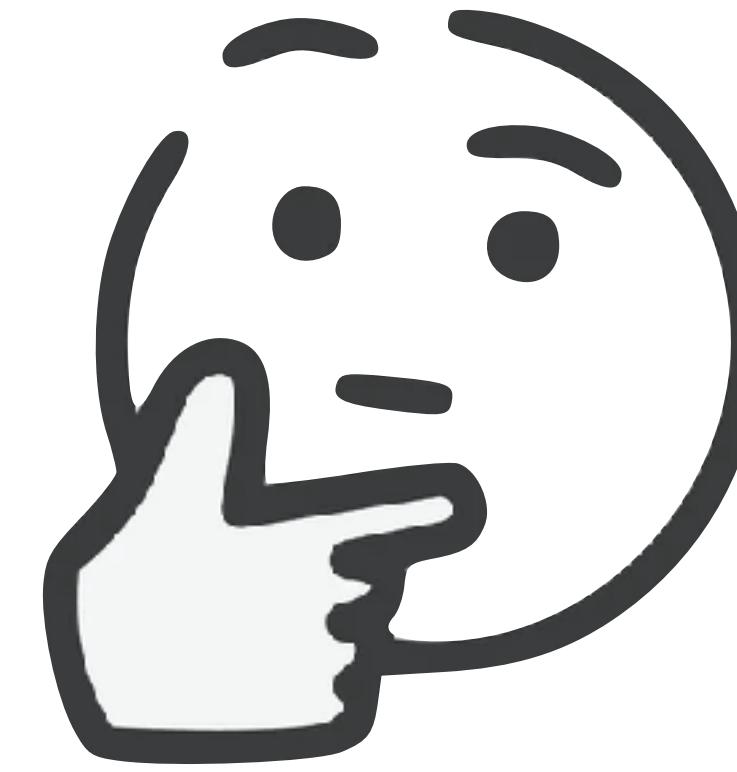
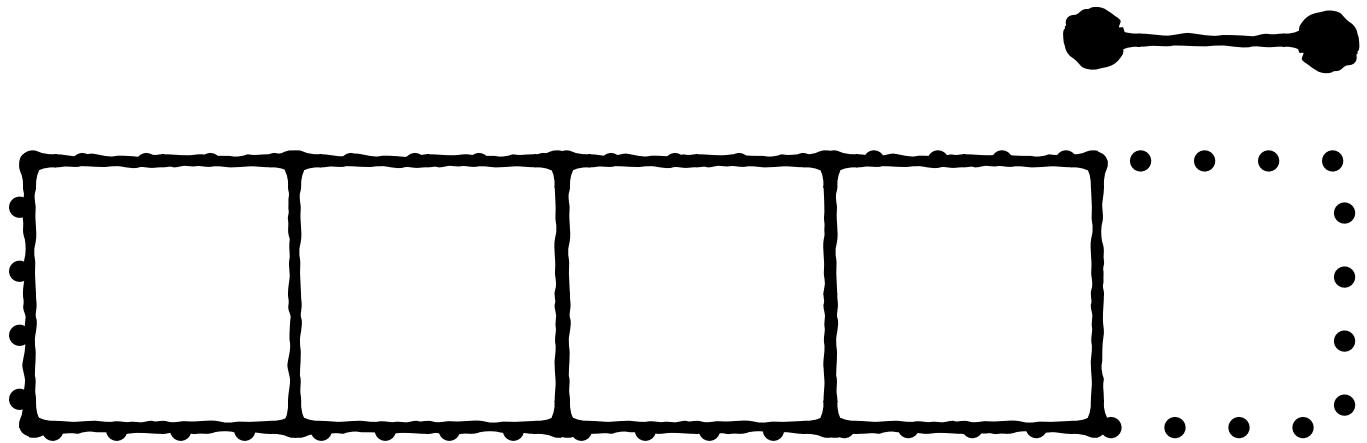


**Could have expanded
by larger factor**

For which growth factor, do we perfectly reuse the space?



For which growth factor, do we perfectly reuse the space?



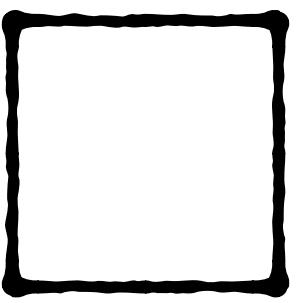
Applicable question across other data structures, e.g., hash tables

Assumptions on memory allocator



Assumptions:

Contiguous allocation



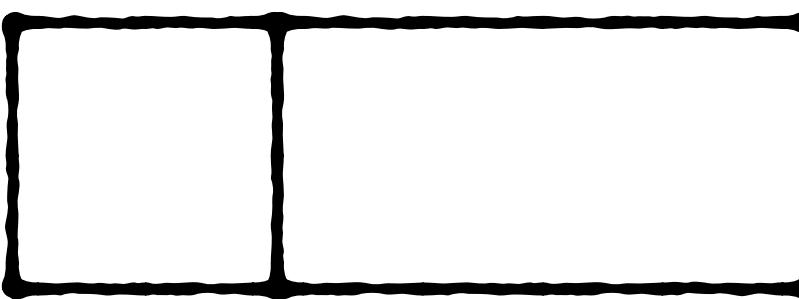
0

Memory addresses



Assumptions:

Contiguous allocation



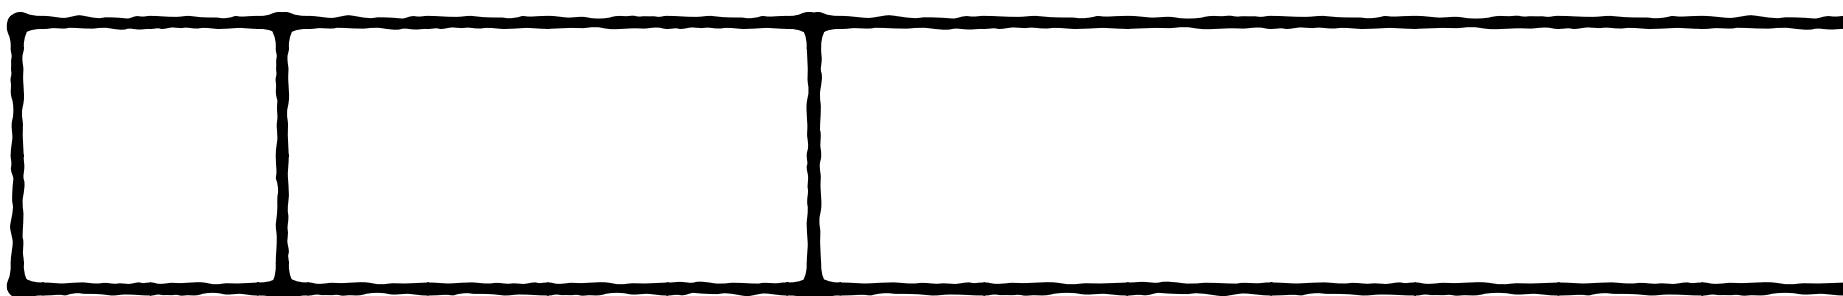
0

Memory addresses



Assumptions:

Contiguous allocation



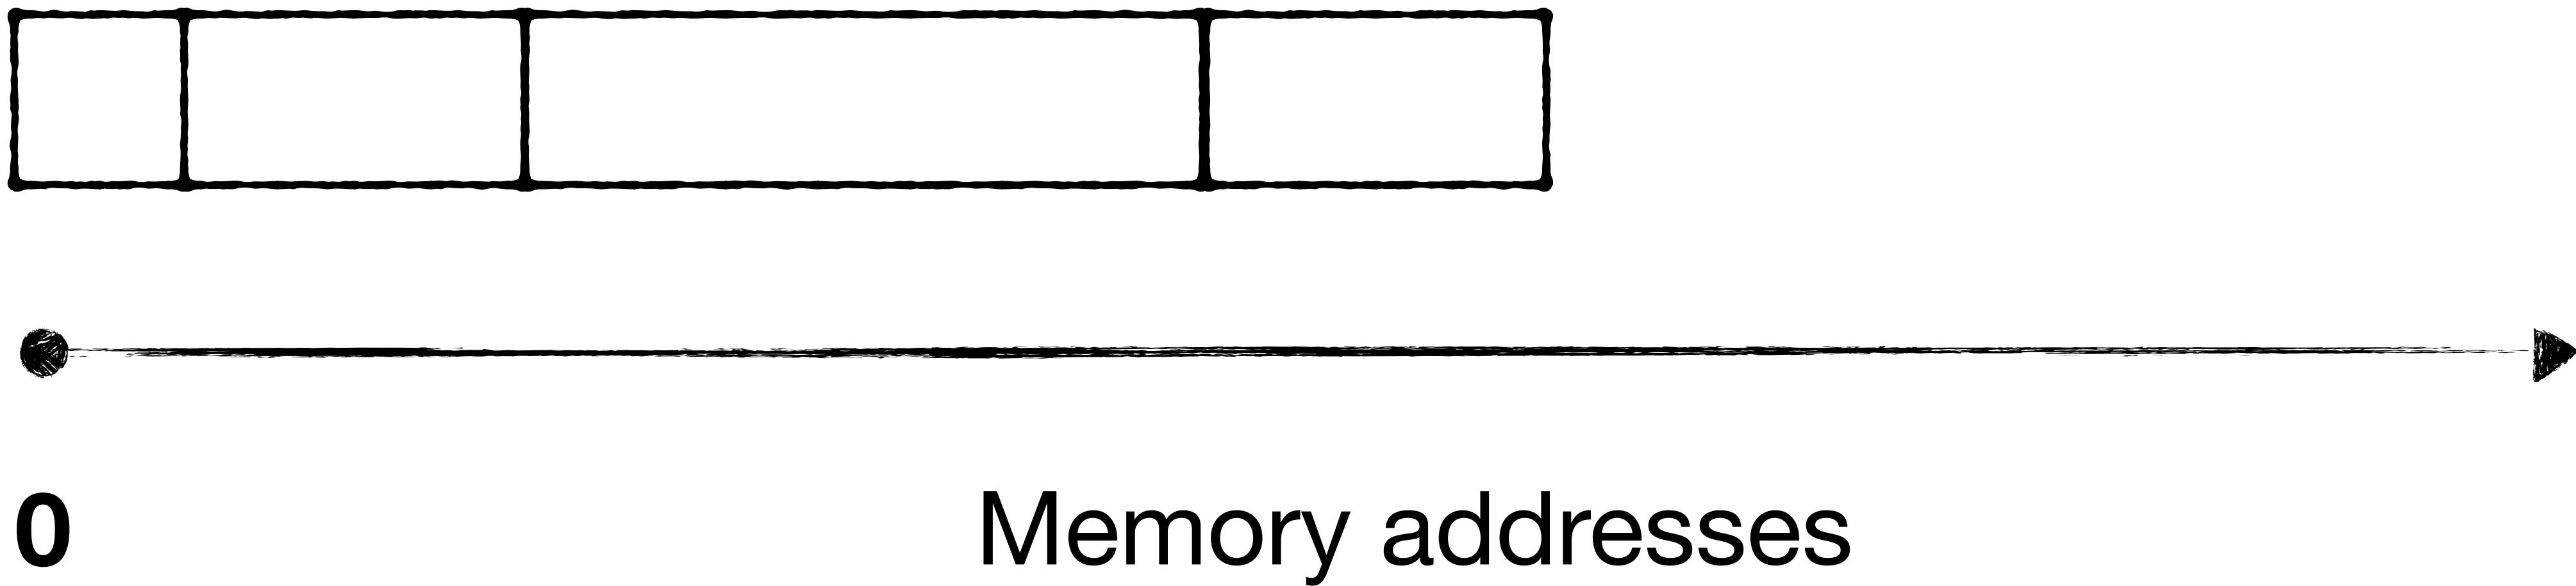
0

Memory addresses



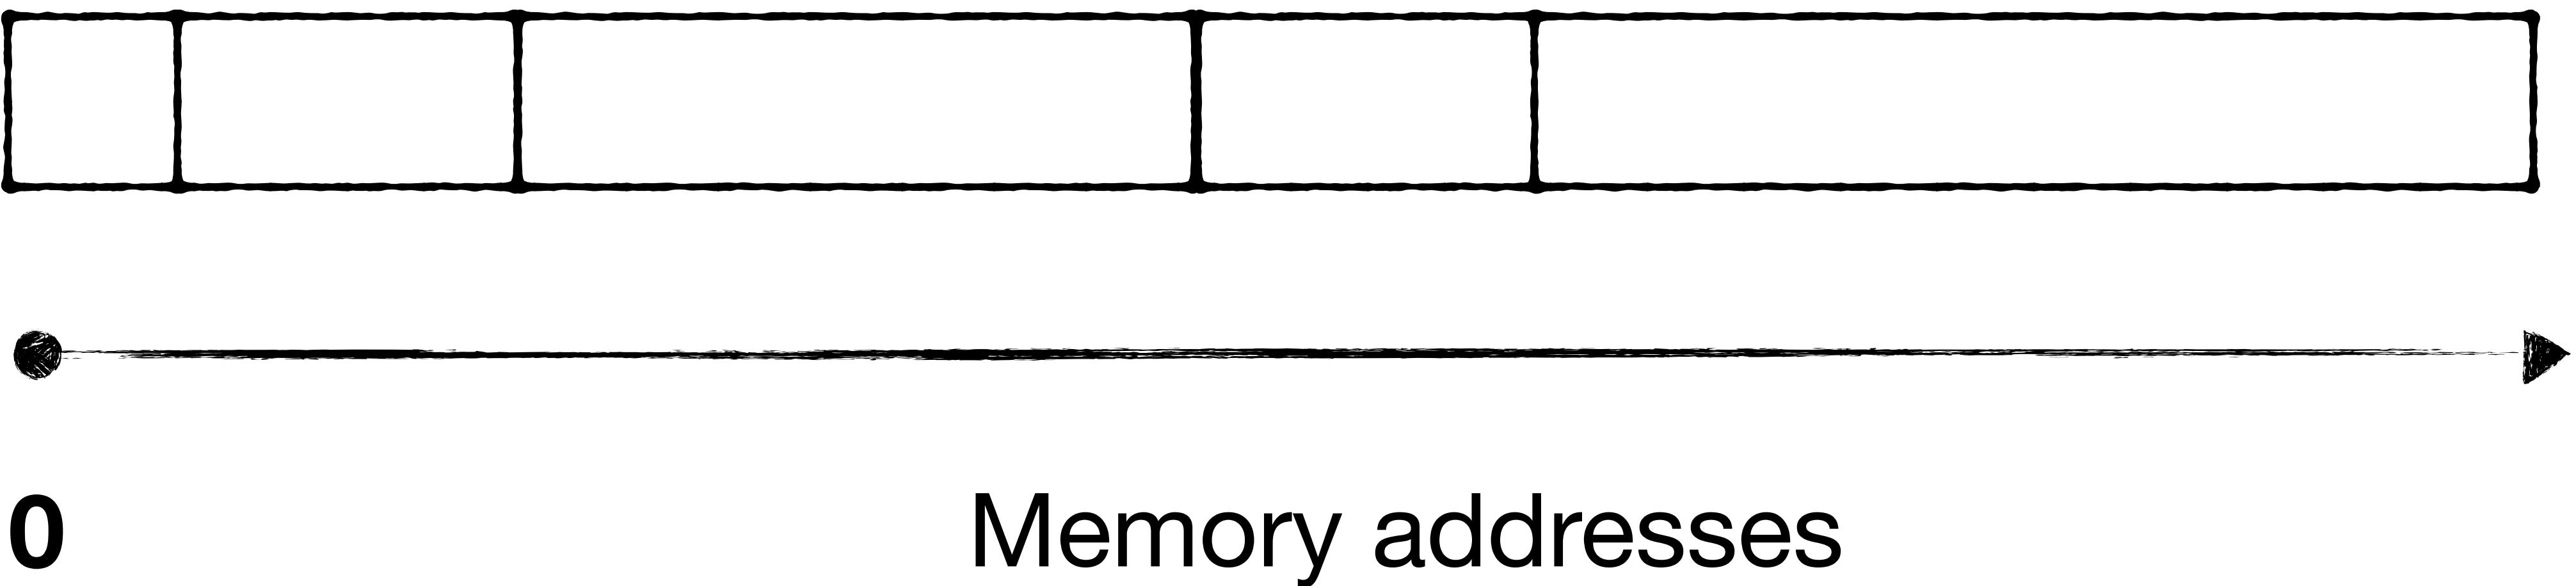
Assumptions:

Contiguous allocation



Assumptions:

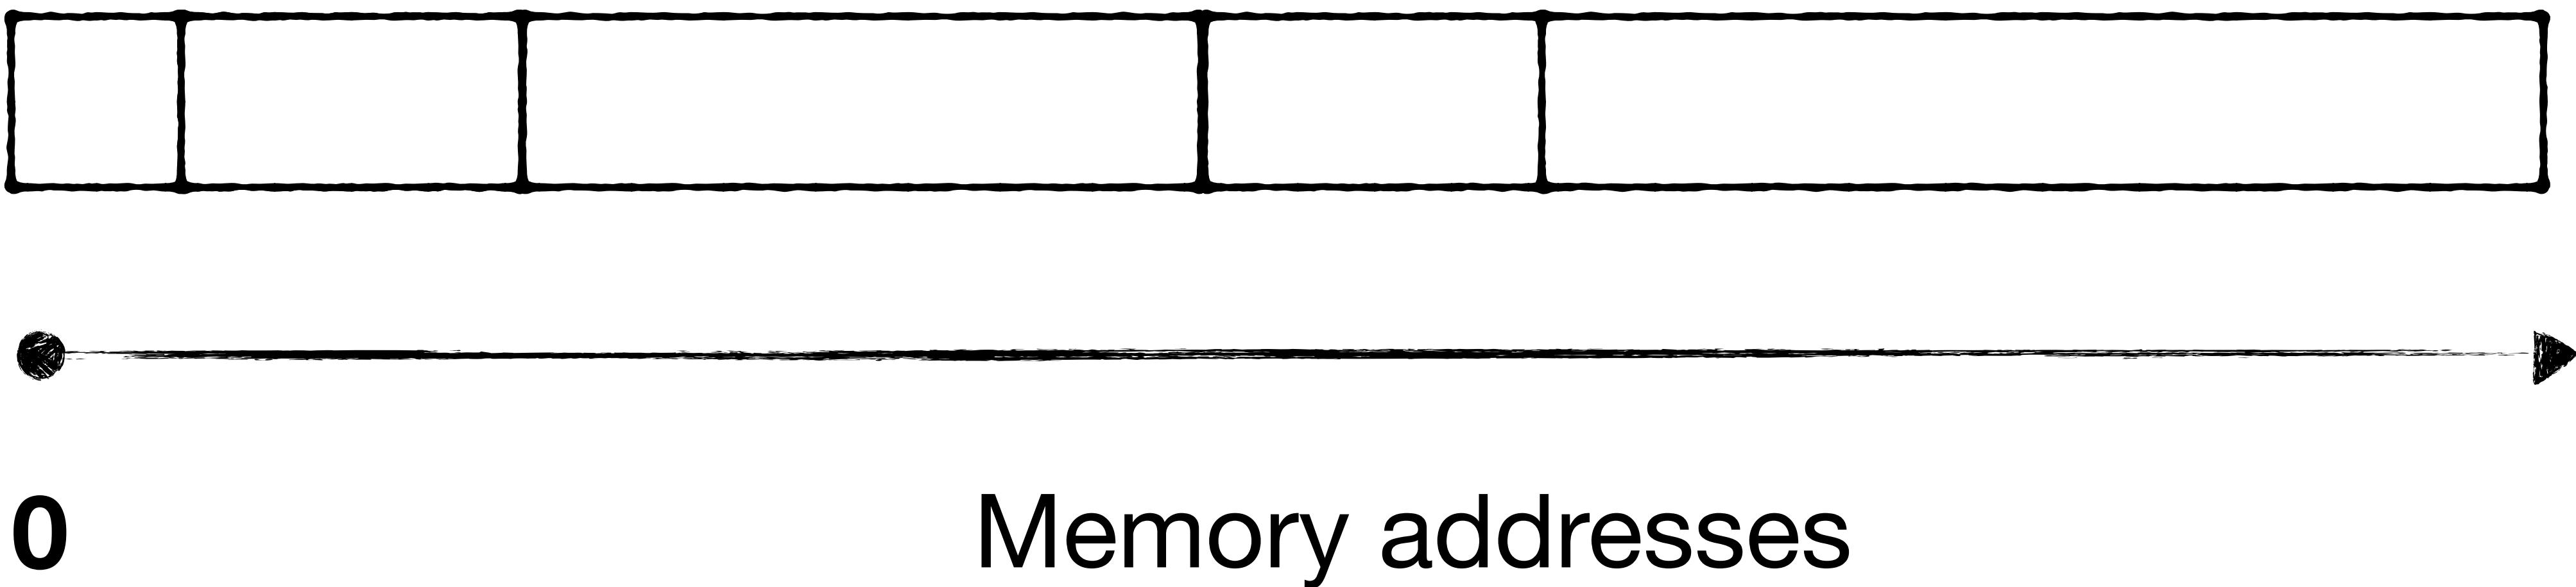
Contiguous allocation



Assumptions:

Contiguous allocation

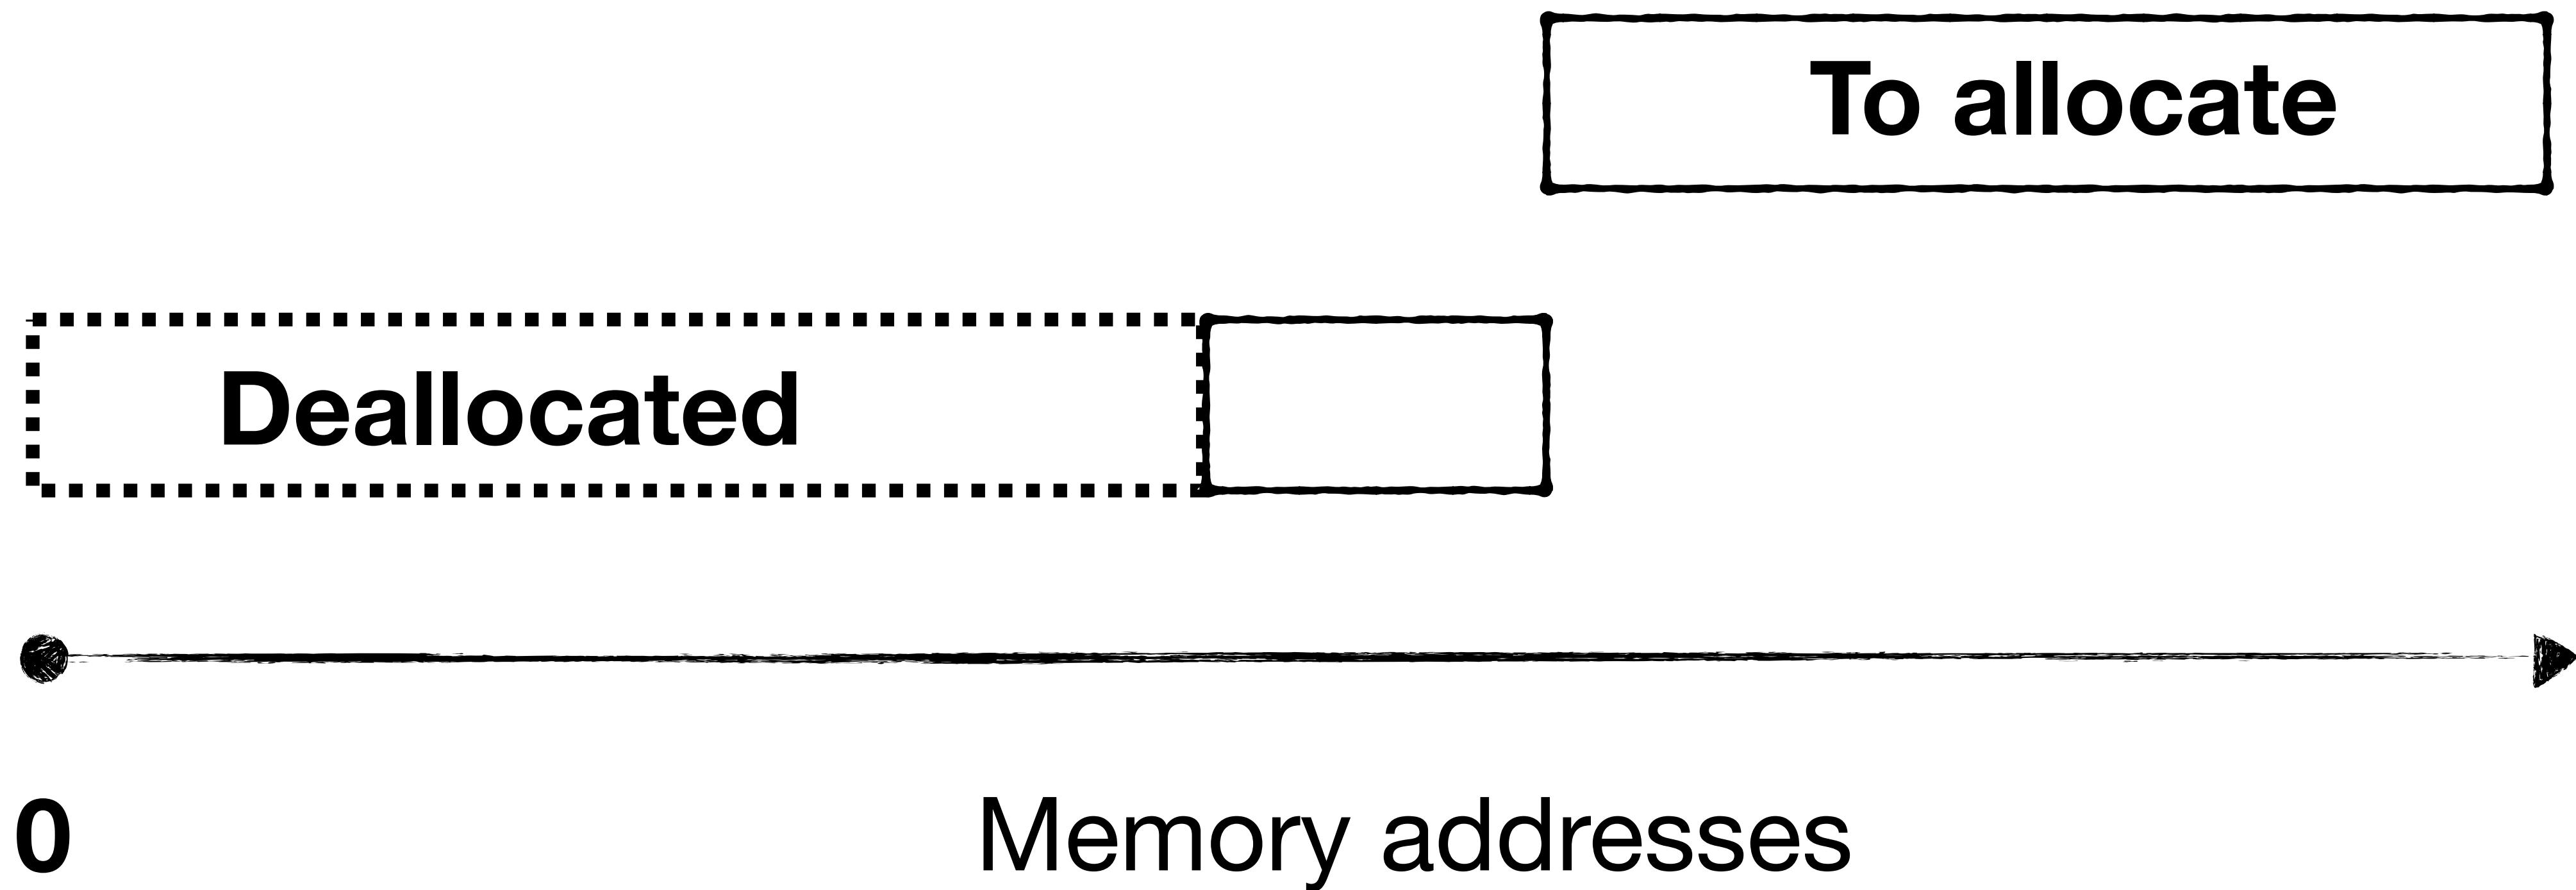
Reuse when possible



Assumptions:

Contiguous allocation

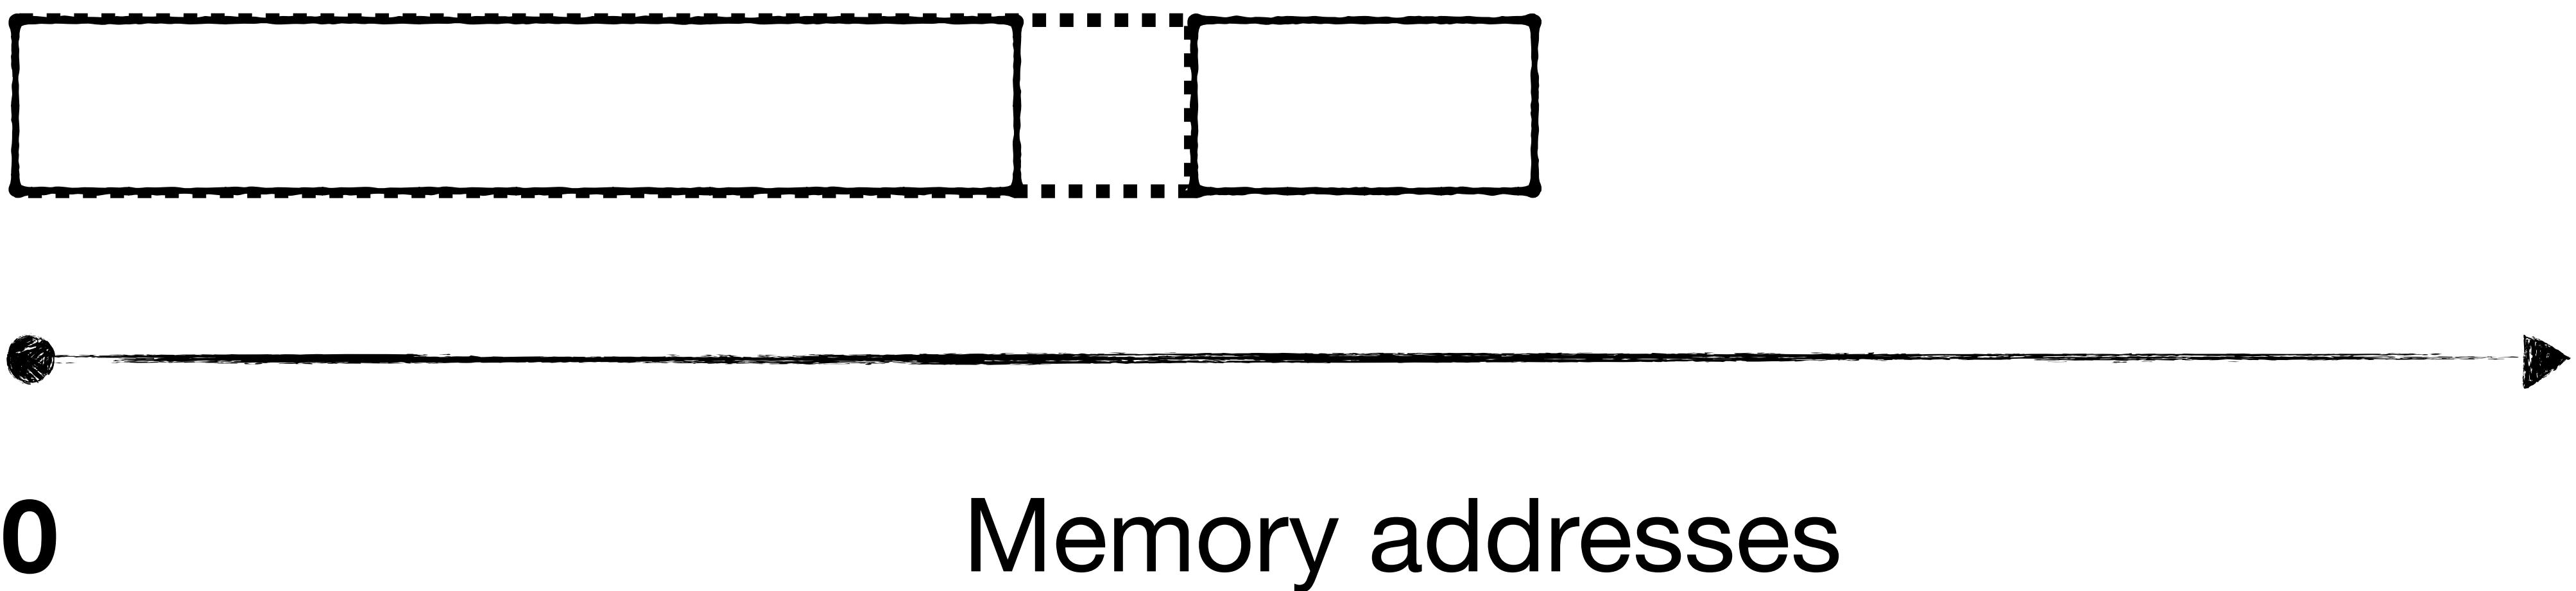
Reuse when possible



Assumptions:

Contiguous allocation

Reuse when possible

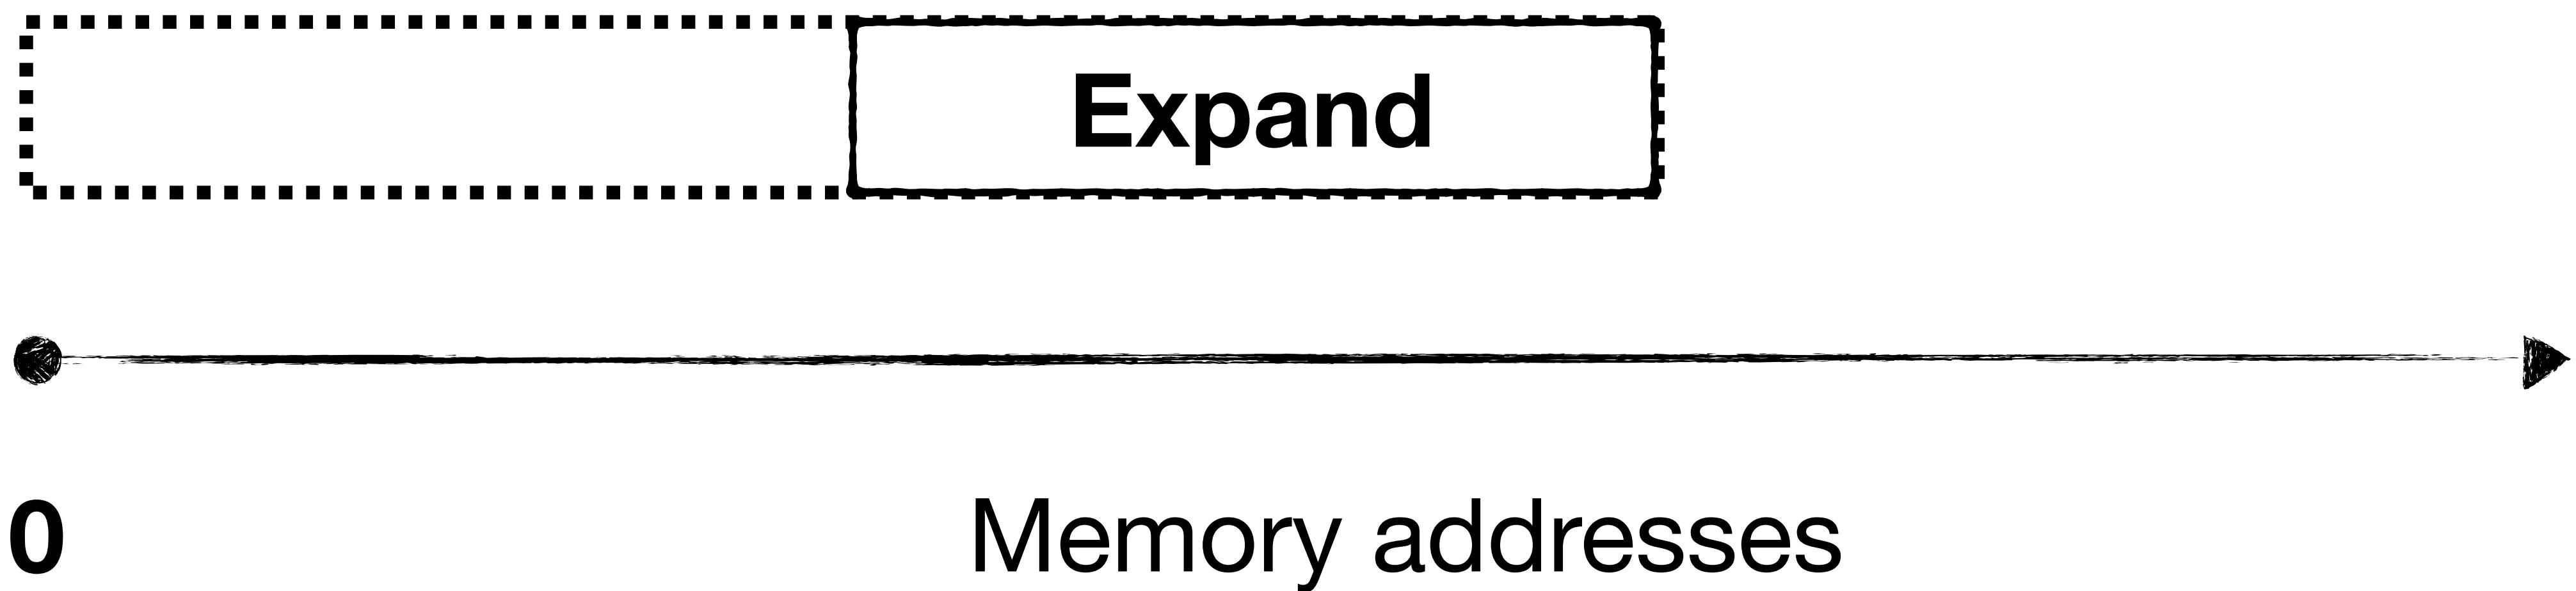


Assumptions:

Contiguous allocation

Reuse when possible

Deallocate as we copy (simplistic)

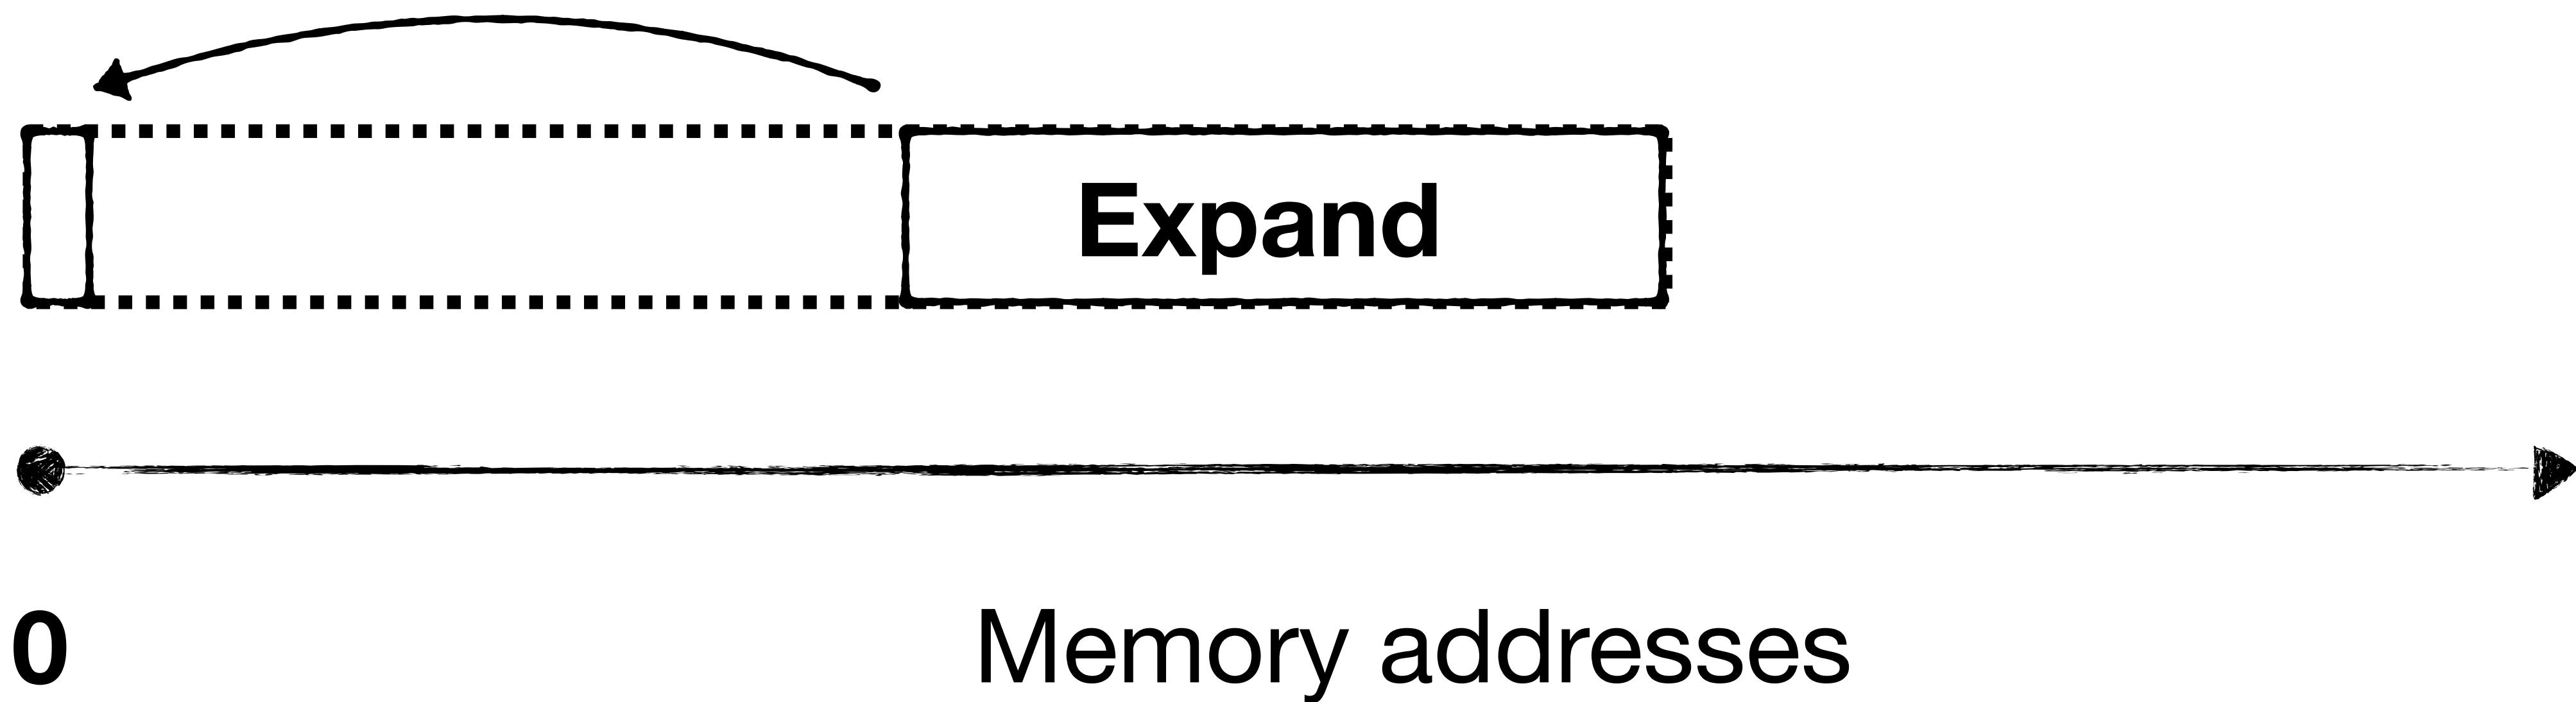


Assumptions:

Contiguous allocation

Reuse when possible

Deallocate as we copy (simplistic)

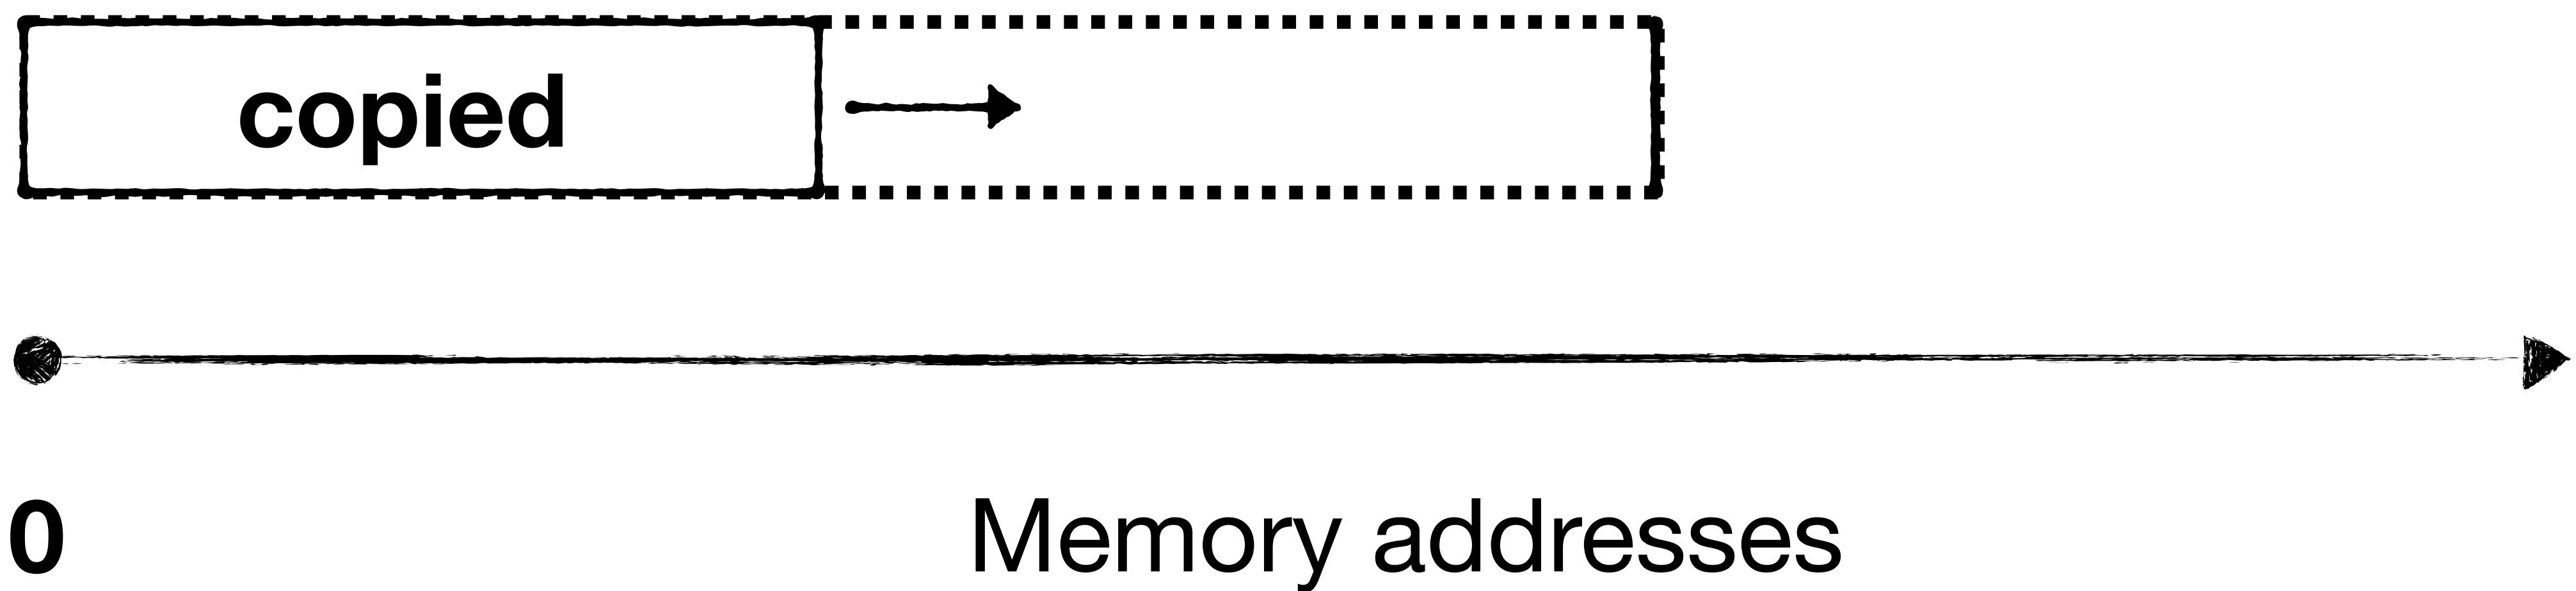


Assumptions:

Contiguous allocation

Reuse when possible

Deallocate as we copy (simplistic)

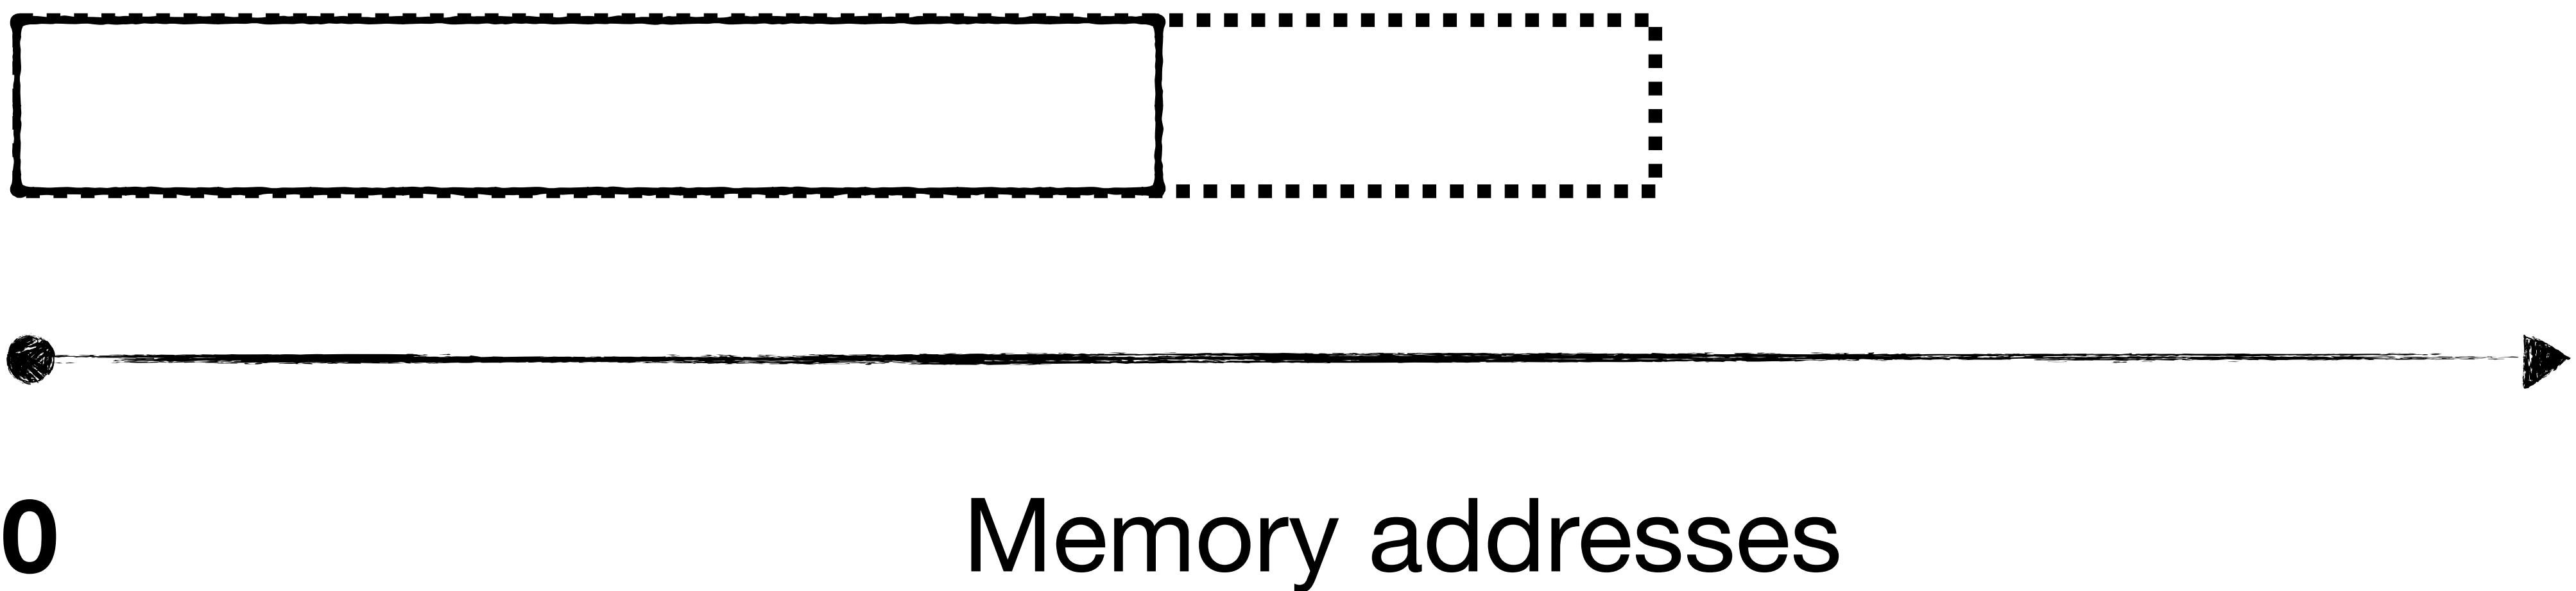


Assumptions:

Contiguous allocation

Reuse when possible

Deallocate as we copy (simplistic)



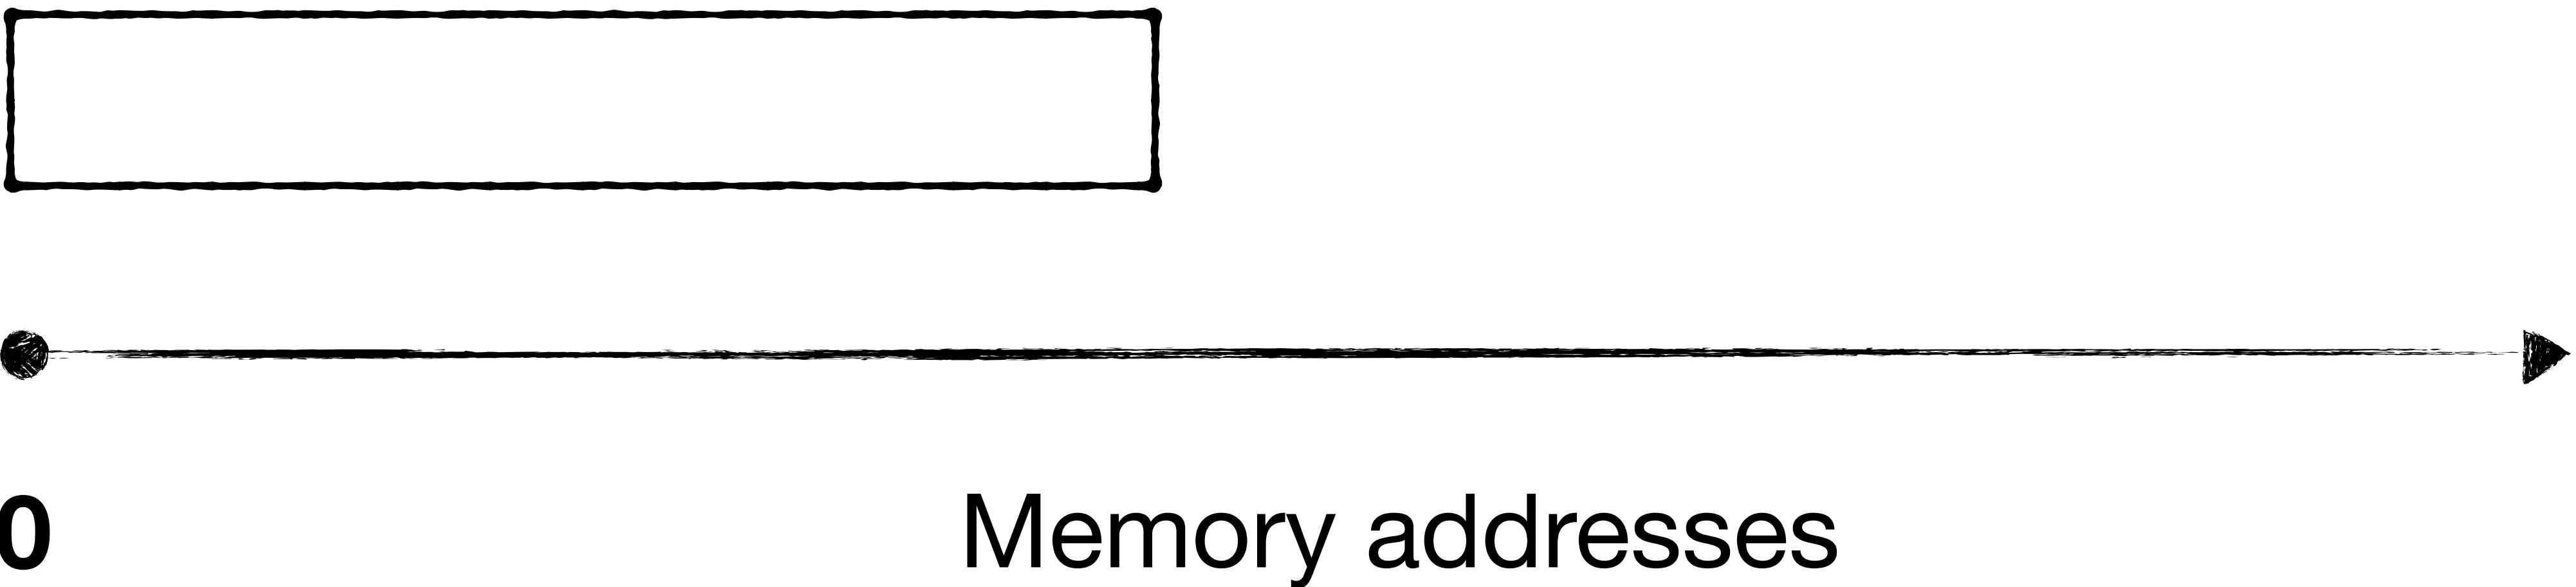
Assumptions:

Contiguous allocation

Reuse when possible

Deallocate as we copy

Can't expand in-place



Assumptions:

Contiguous allocation

Reuse when possible

Deallocate as we copy

Can't expand in-place



0

Memory addresses

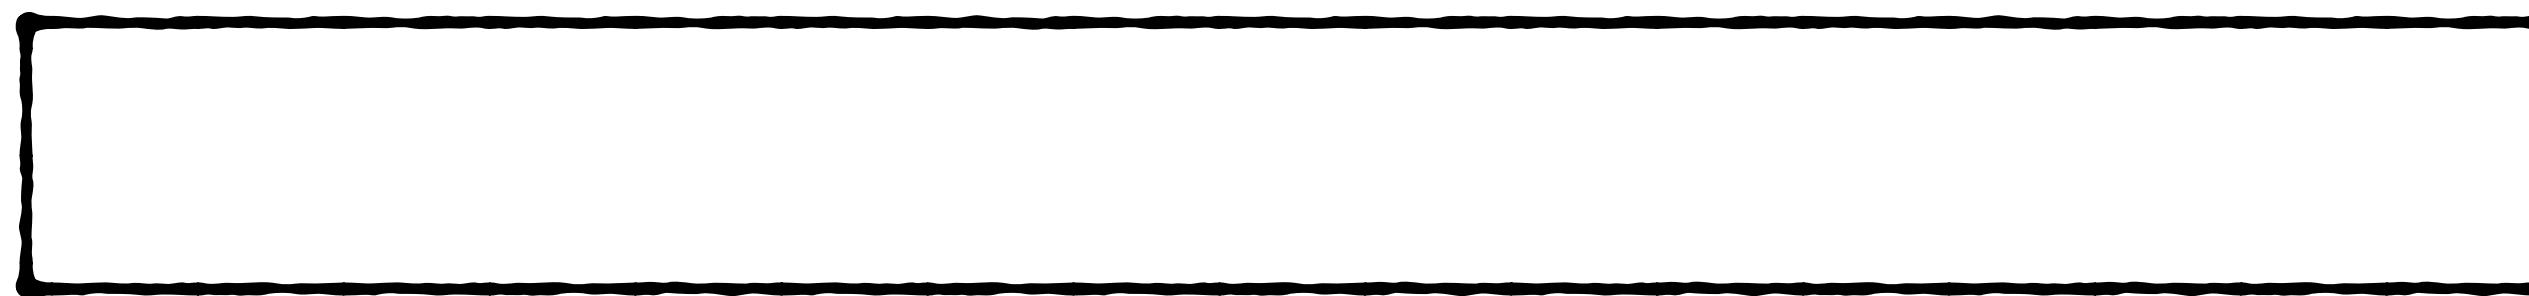
Assumptions:

Contiguous allocation

Reuse when possible

Deallocate as we copy

Can't expand in-place



0

Memory addresses

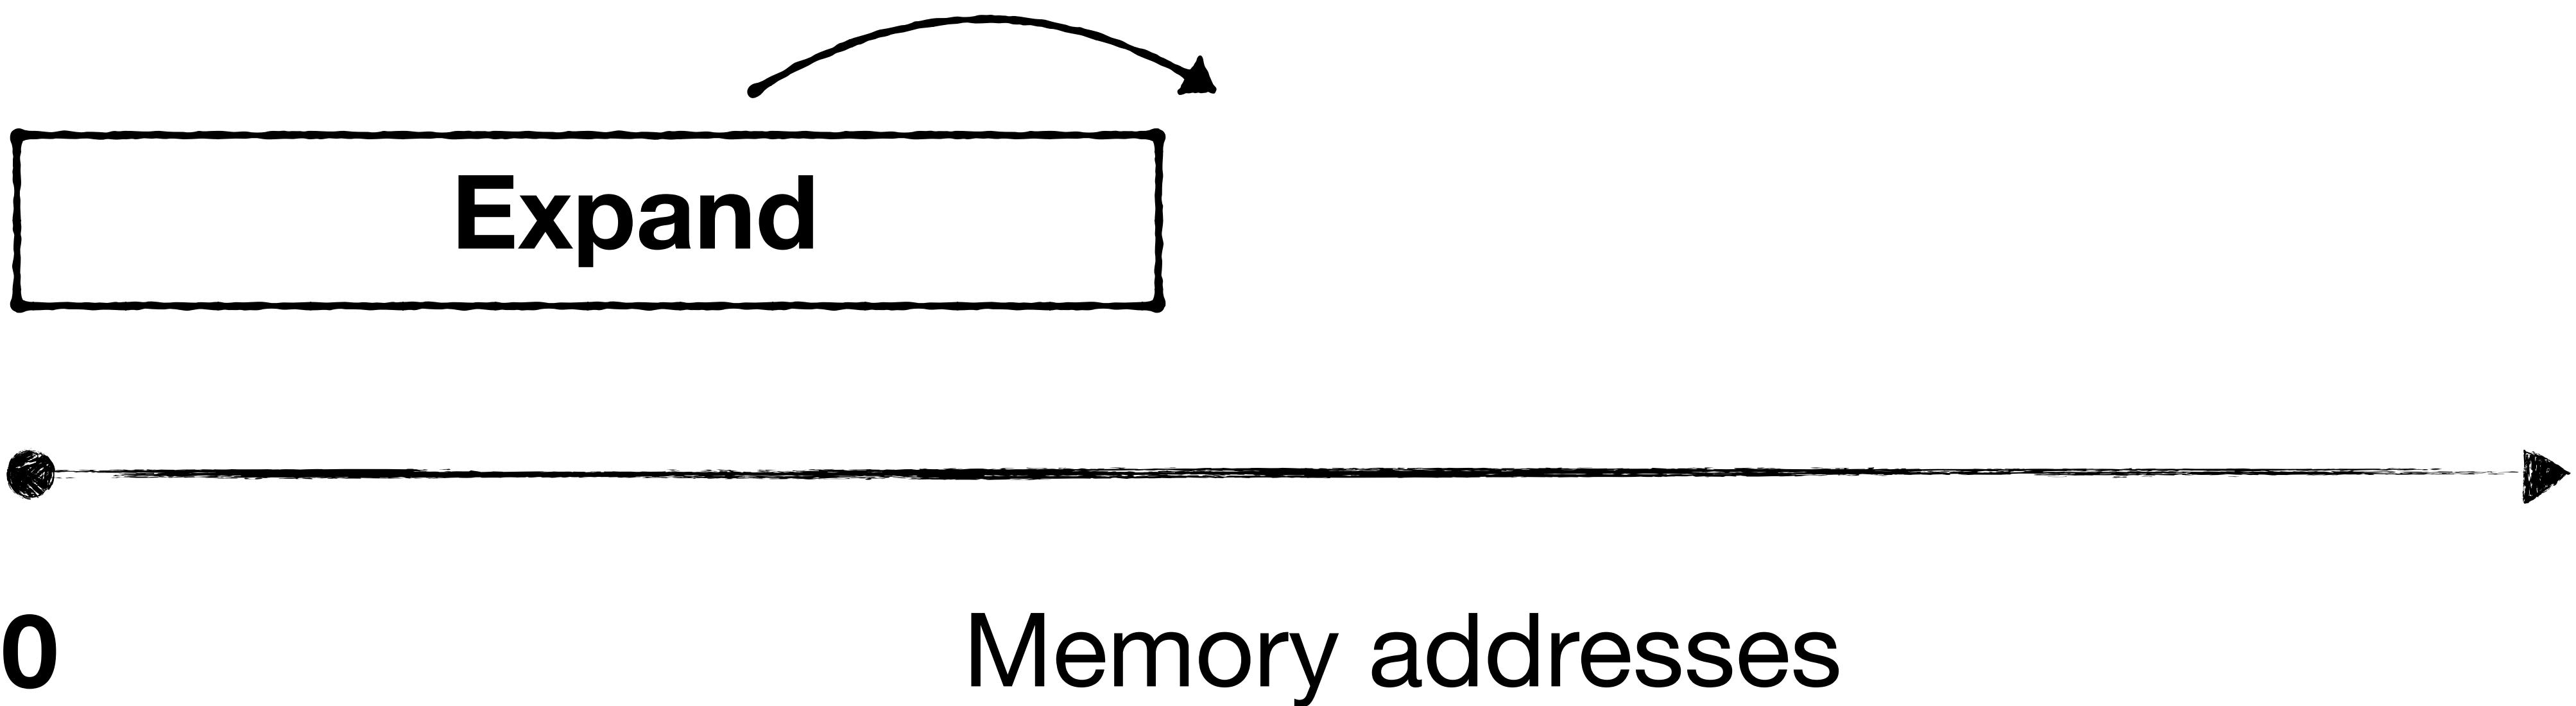
Assumptions:

Contiguous allocation

Reuse when possible

Deallocate as we copy

Can't expand in-place



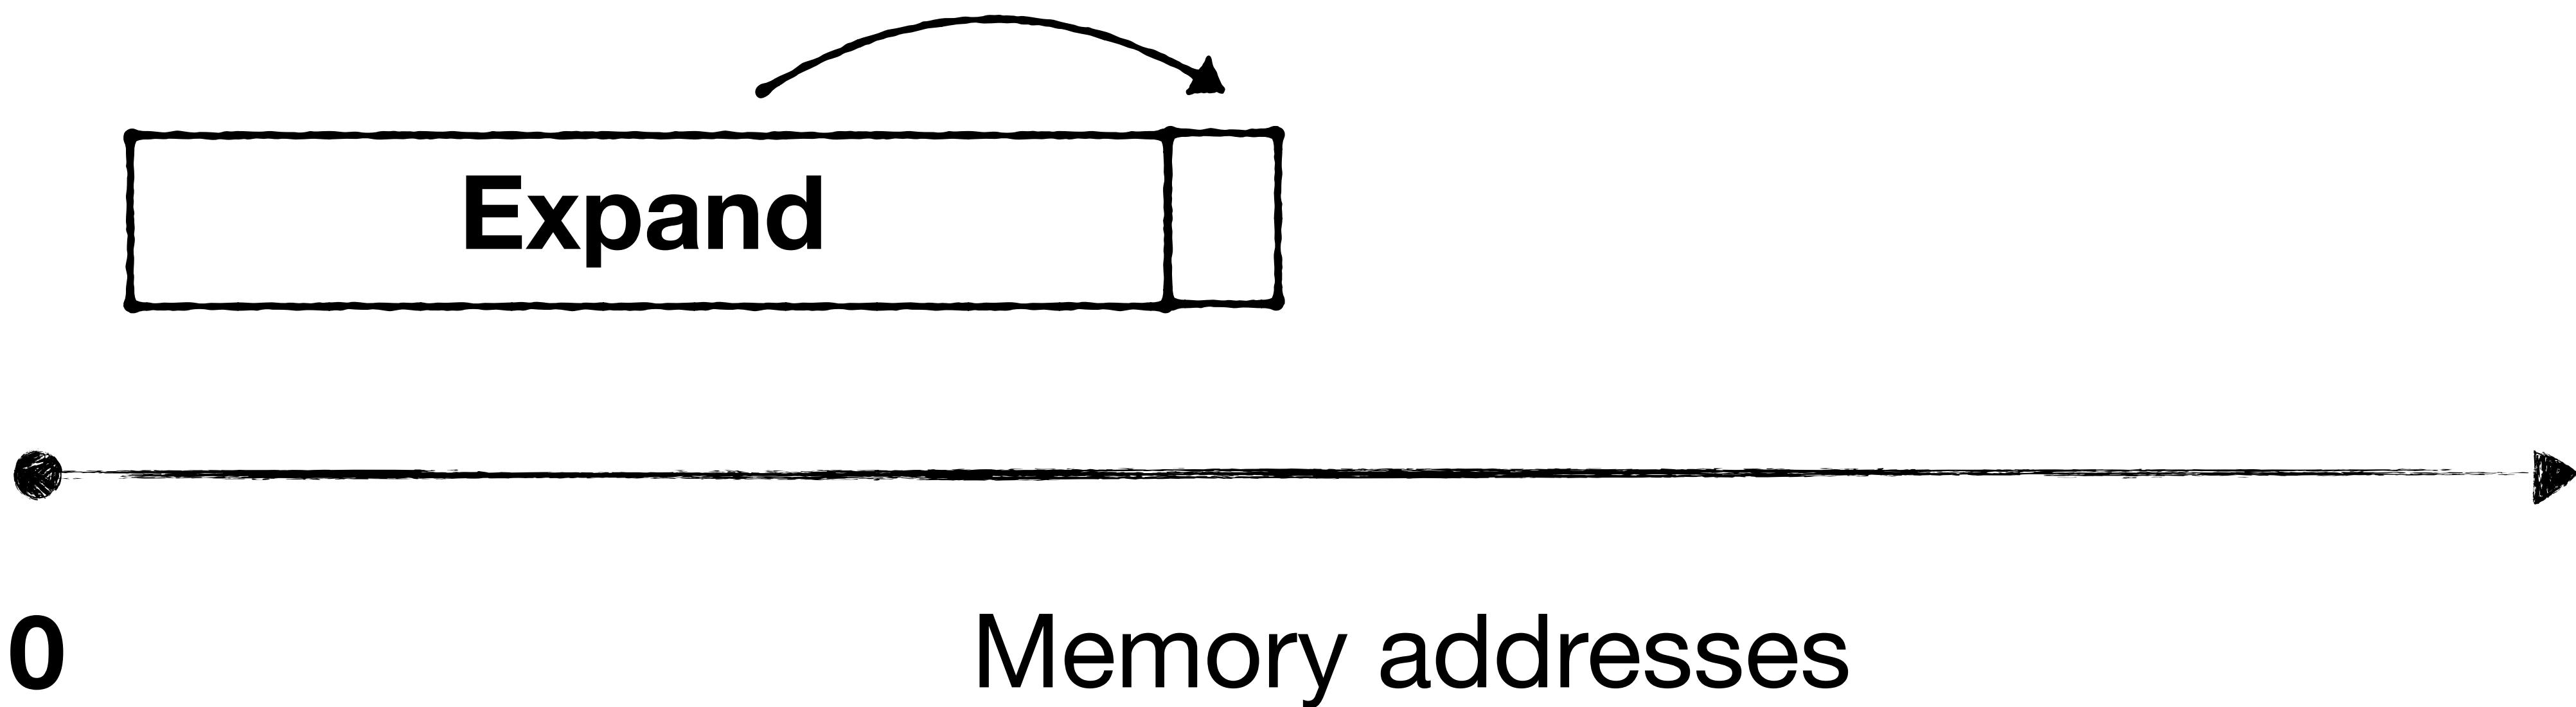
Assumptions:

Contiguous allocation

Reuse when possible

Deallocate as we copy

Can't expand in-place



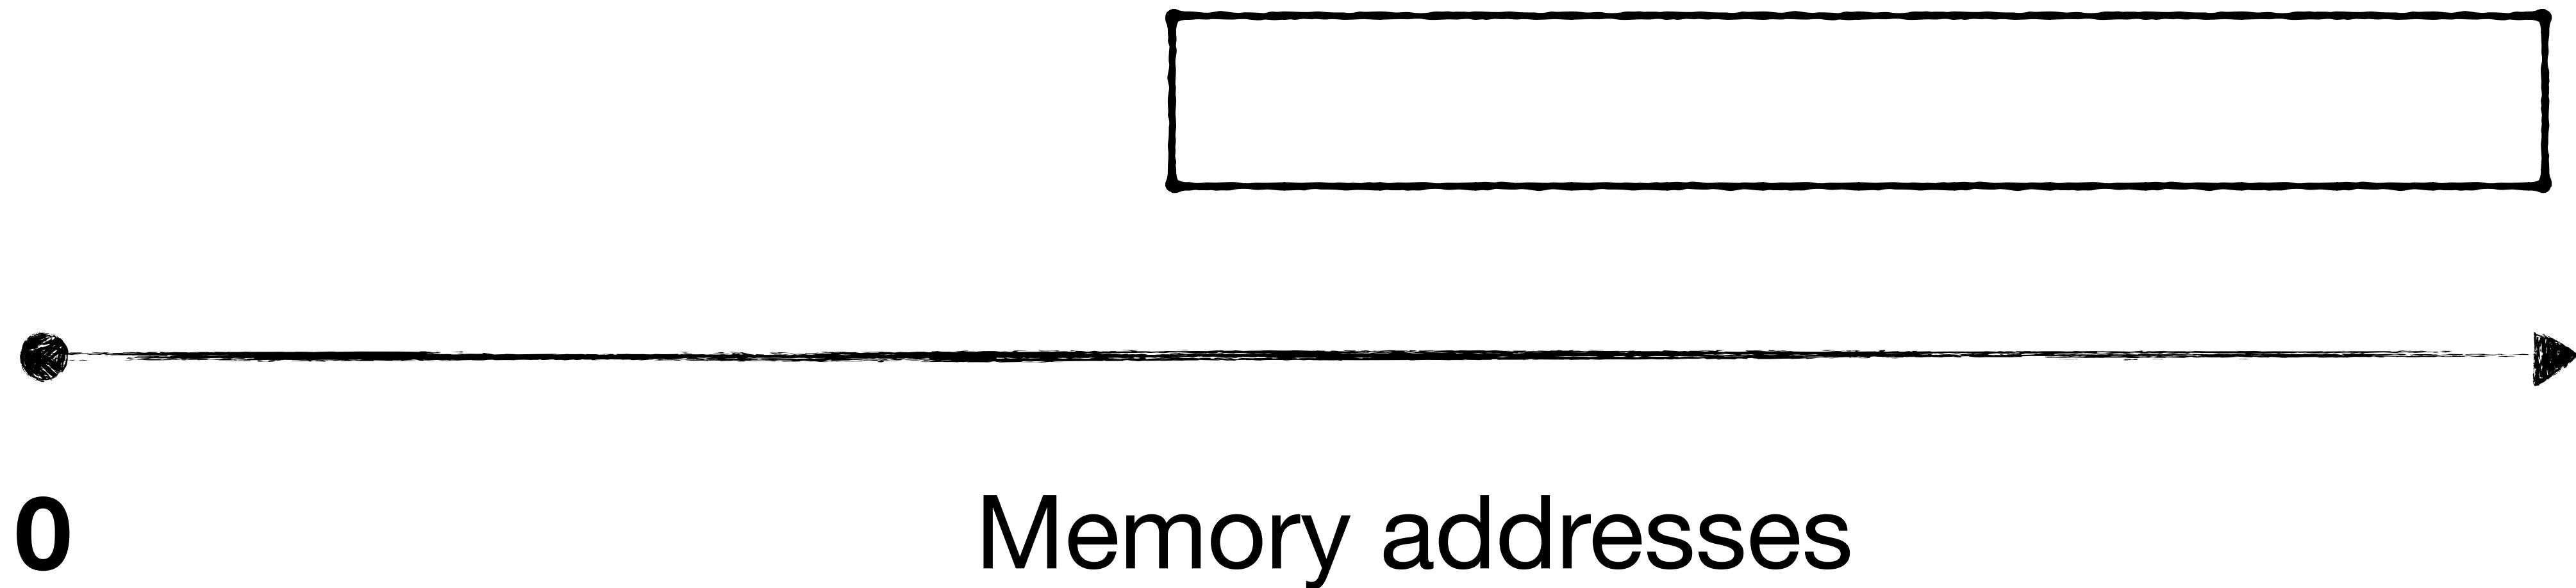
Assumptions:

Contiguous allocation

Reuse when possible

Deallocate as we copy

Can't expand in-place



For which growth factor, do we perfectly reuse the space?

For which growth factor, do we perfectly reuse the space?

$$\mathbf{Size}_{i-2} + \mathbf{Size}_{i-1} = \mathbf{Size}_i$$

For which growth factor, do we perfectly reuse the space?

$$\mathbf{Size}_{i-2} + \mathbf{Size}_{i-1} = \mathbf{Size}_i$$

Subject to:

$$\frac{\mathbf{Size}_{i-1}}{\mathbf{Size}_{i-2}} = \frac{\mathbf{Size}_i}{\mathbf{Size}_{i-1}} = G$$

For which growth factor, do we perfectly reuse the space?

$$\text{Size}_{i-2} + \text{Size}_{i-1} = \text{Size}_i$$

Subject to:

$$\frac{\text{Size}_{i-1}}{\text{Size}_{i-2}} = \frac{\text{Size}_i}{\text{Size}_{i-1}} = G$$

Clever ideas?



Fibonacci Series

1 1 2 3 5 8 13 21

$$\text{Size}_{i-2} + \text{Size}_{i-1} = \text{Size}_i$$

$$\frac{\text{Size}_{i-1}}{\text{Size}_{i-2}} = \frac{\text{Size}_i}{\text{Size}_{i-1}} = G$$

Fibonacci Series

1 1 2 3 5 8 13 21



1170 - 1250

Italy

$$\text{Size}_{i-2} + \text{Size}_{i-1} = \text{Size}_i$$

$$\frac{\text{Size}_{i-1}}{\text{Size}_{i-2}} = \frac{\text{Size}_i}{\text{Size}_{i-1}} = G$$

Fibonacci Series

1 + 1 = 2 3 5 8 13 21

$$\text{Size}_{i-2} + \text{Size}_{i-1} = \text{Size}_i$$

$$\frac{\text{Size}_{i-1}}{\text{Size}_{i-2}} = \frac{\text{Size}_i}{\text{Size}_{i-1}} = G$$

Fibonacci Series

1 1 + 2 = 3 5 8 13 21

$$\text{Size}_{i-2} + \text{Size}_{i-1} = \text{Size}_i$$

$$\frac{\text{Size}_{i-1}}{\text{Size}_{i-2}} = \frac{\text{Size}_i}{\text{Size}_{i-1}} = G$$

Fibonacci Series

1 1 2 + 3 = 5 8 13 21

$$\text{Size}_{i-2} + \text{Size}_{i-1} = \text{Size}_i$$

$$\frac{\text{Size}_{i-1}}{\text{Size}_{i-2}} = \frac{\text{Size}_i}{\text{Size}_{i-1}} = G$$

Fibonacci Series

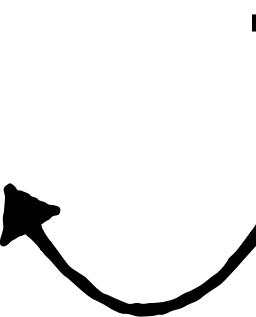
1 1 2 + 3 = 5 8 13 21

Satisfies this: \rightarrow **Size _{i-2}** + **Size _{i-1}** = **Size _i**

$$\frac{\text{Size}_{i-1}}{\text{Size}_{i-2}} = \frac{\text{Size}_i}{\text{Size}_{i-1}} = G$$

Fibonacci Series

1 1 2 3 5 8 13 21



1

$$\text{Size}_{i-2} + \text{Size}_{i-1} = \text{Size}_i$$

$$\frac{\text{Size}_{i-1}}{\text{Size}_{i-2}} = \frac{\text{Size}_i}{\text{Size}_{i-1}} = G$$

Fibonacci Series

1 1 2 3 5 8 13 21



2

$$\text{Size}_{i-2} + \text{Size}_{i-1} = \text{Size}_i$$

$$\frac{\text{Size}_{i-1}}{\text{Size}_{i-2}} = \frac{\text{Size}_i}{\text{Size}_{i-1}} = G$$

Fibonacci Series

1 1 2 3 5 8 13 21
 ↗
 1.5

$$\text{Size}_{i-2} + \text{Size}_{i-1} = \text{Size}_i$$

$$\frac{\text{Size}_{i-1}}{\text{Size}_{i-2}} = \frac{\text{Size}_i}{\text{Size}_{i-1}} = G$$

Fibonacci Series

1 1 2 3 5 8 13 21
 ↑
1.666

$$\text{Size}_{i-2} + \text{Size}_{i-1} = \text{Size}_i$$

$$\frac{\text{Size}_{i-1}}{\text{Size}_{i-2}} = \frac{\text{Size}_i}{\text{Size}_{i-1}} = G$$

Fibonacci Series

1 1 2 3 5 8 13 21



1.6

$$\text{Size}_{i-2} + \text{Size}_{i-1} = \text{Size}_i$$

$$\frac{\text{Size}_{i-1}}{\text{Size}_{i-2}} = \frac{\text{Size}_i}{\text{Size}_{i-1}} = G$$

Fibonacci Series

1 1 2 3 5 8 13 21

1.625



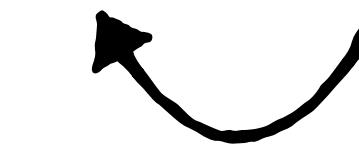
$$\text{Size}_{i-2} + \text{Size}_{i-1} = \text{Size}_i$$

$$\frac{\text{Size}_{i-1}}{\text{Size}_{i-2}} = \frac{\text{Size}_i}{\text{Size}_{i-1}} = G$$

Fibonacci Series

1 1 2 3 5 8 13 21

1.625



$$\text{Size}_{i-2} + \text{Size}_{i-1} = \text{Size}_i$$

$$\frac{\text{Size}_{i-1}}{\text{Size}_{i-2}} = \frac{\text{Size}_i}{\text{Size}_{i-1}} = G$$

Fibonacci Series

1 1 2 3 5 8 13 21



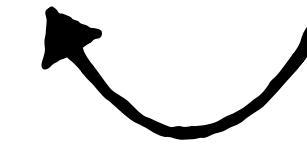
$$\text{Size}_{i-2} + \text{Size}_{i-1} = \text{Size}_i$$

$$\frac{\text{Size}_{i-1}}{\text{Size}_{i-2}} = \frac{\text{Size}_i}{\text{Size}_{i-1}} = G$$

Fibonacci Series

1 1 2 3 5 8 13 21 34 55 89 144

1.618

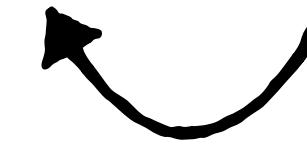


$$\text{Size}_{i-2} + \text{Size}_{i-1} = \text{Size}_i$$

$$\frac{\text{Size}_{i-1}}{\text{Size}_{i-2}} = \frac{\text{Size}_i}{\text{Size}_{i-1}} = G$$

Fibonacci Series

1 1 2 3 5 8 13 21 34 55 89 144

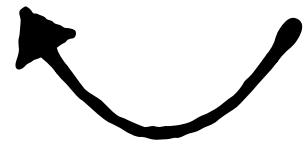


1.618

Ratio converges to the “Golden Ratio” ϕ

Fibonacci Series

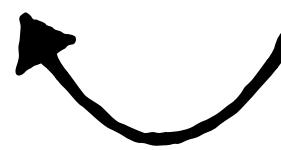
1 1 2 3 5 8 13 21 34 55 89 144



Ratio converges to the “Golden Ratio” $\phi = 1.618033988749\dots$

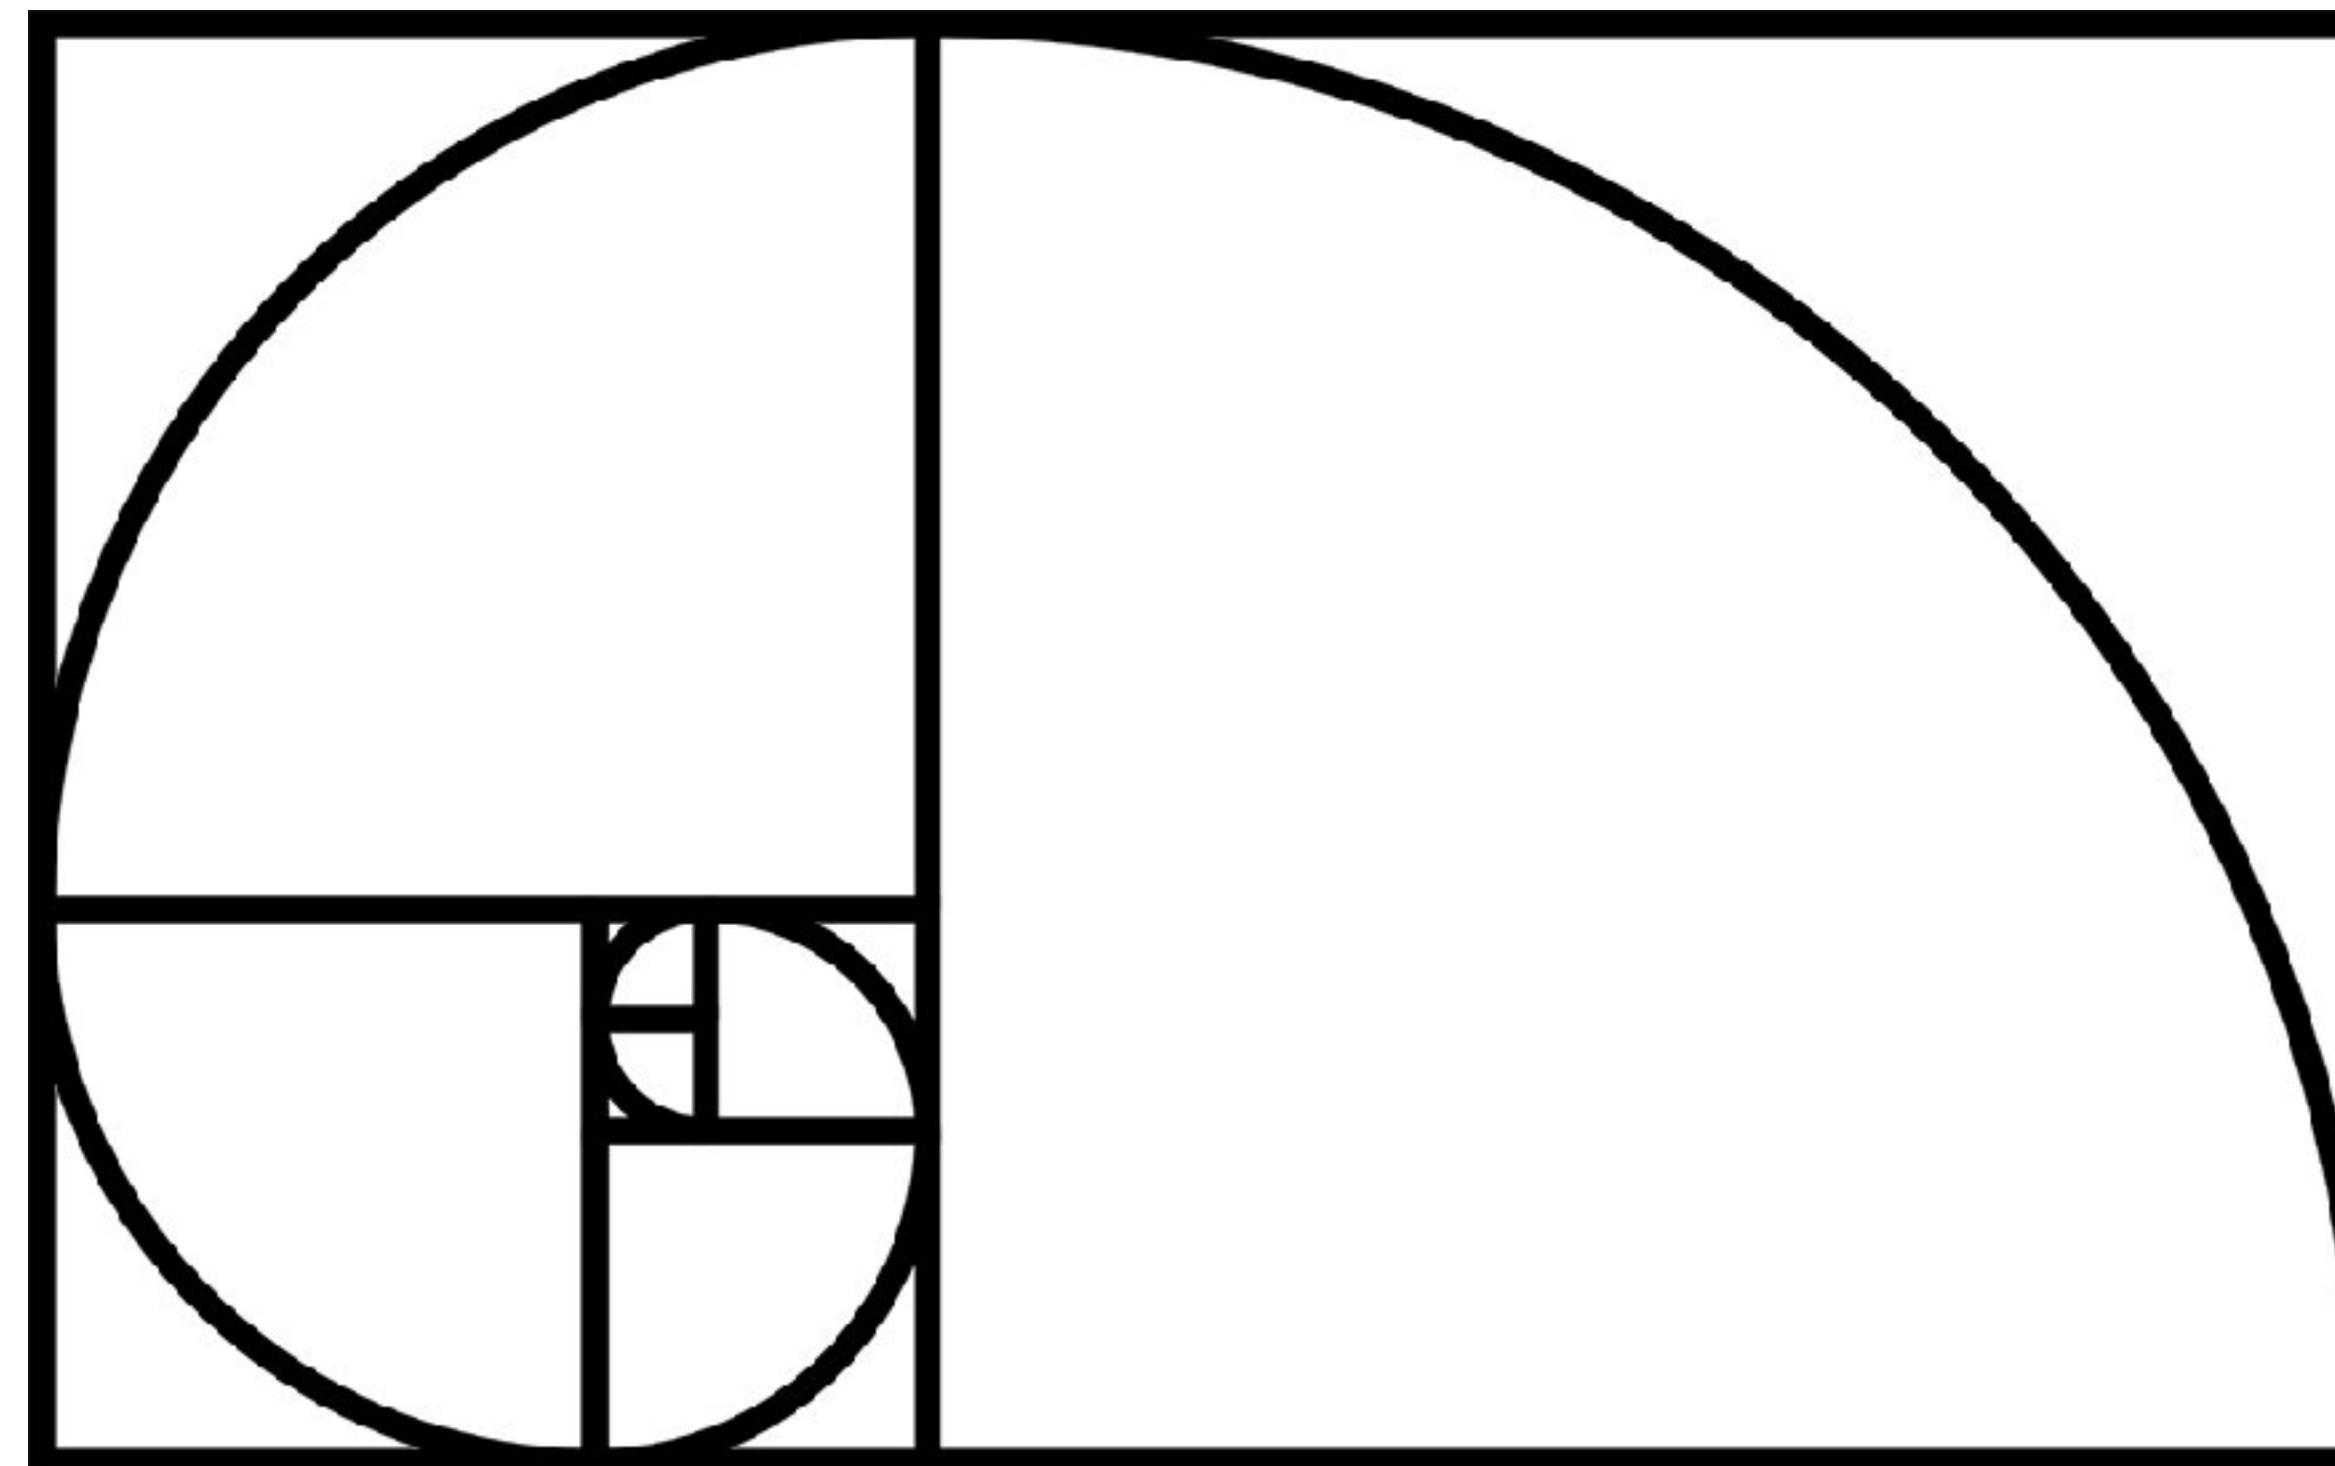
Fibonacci Series

1 1 2 3 5 8 13 21 34 55 89 144



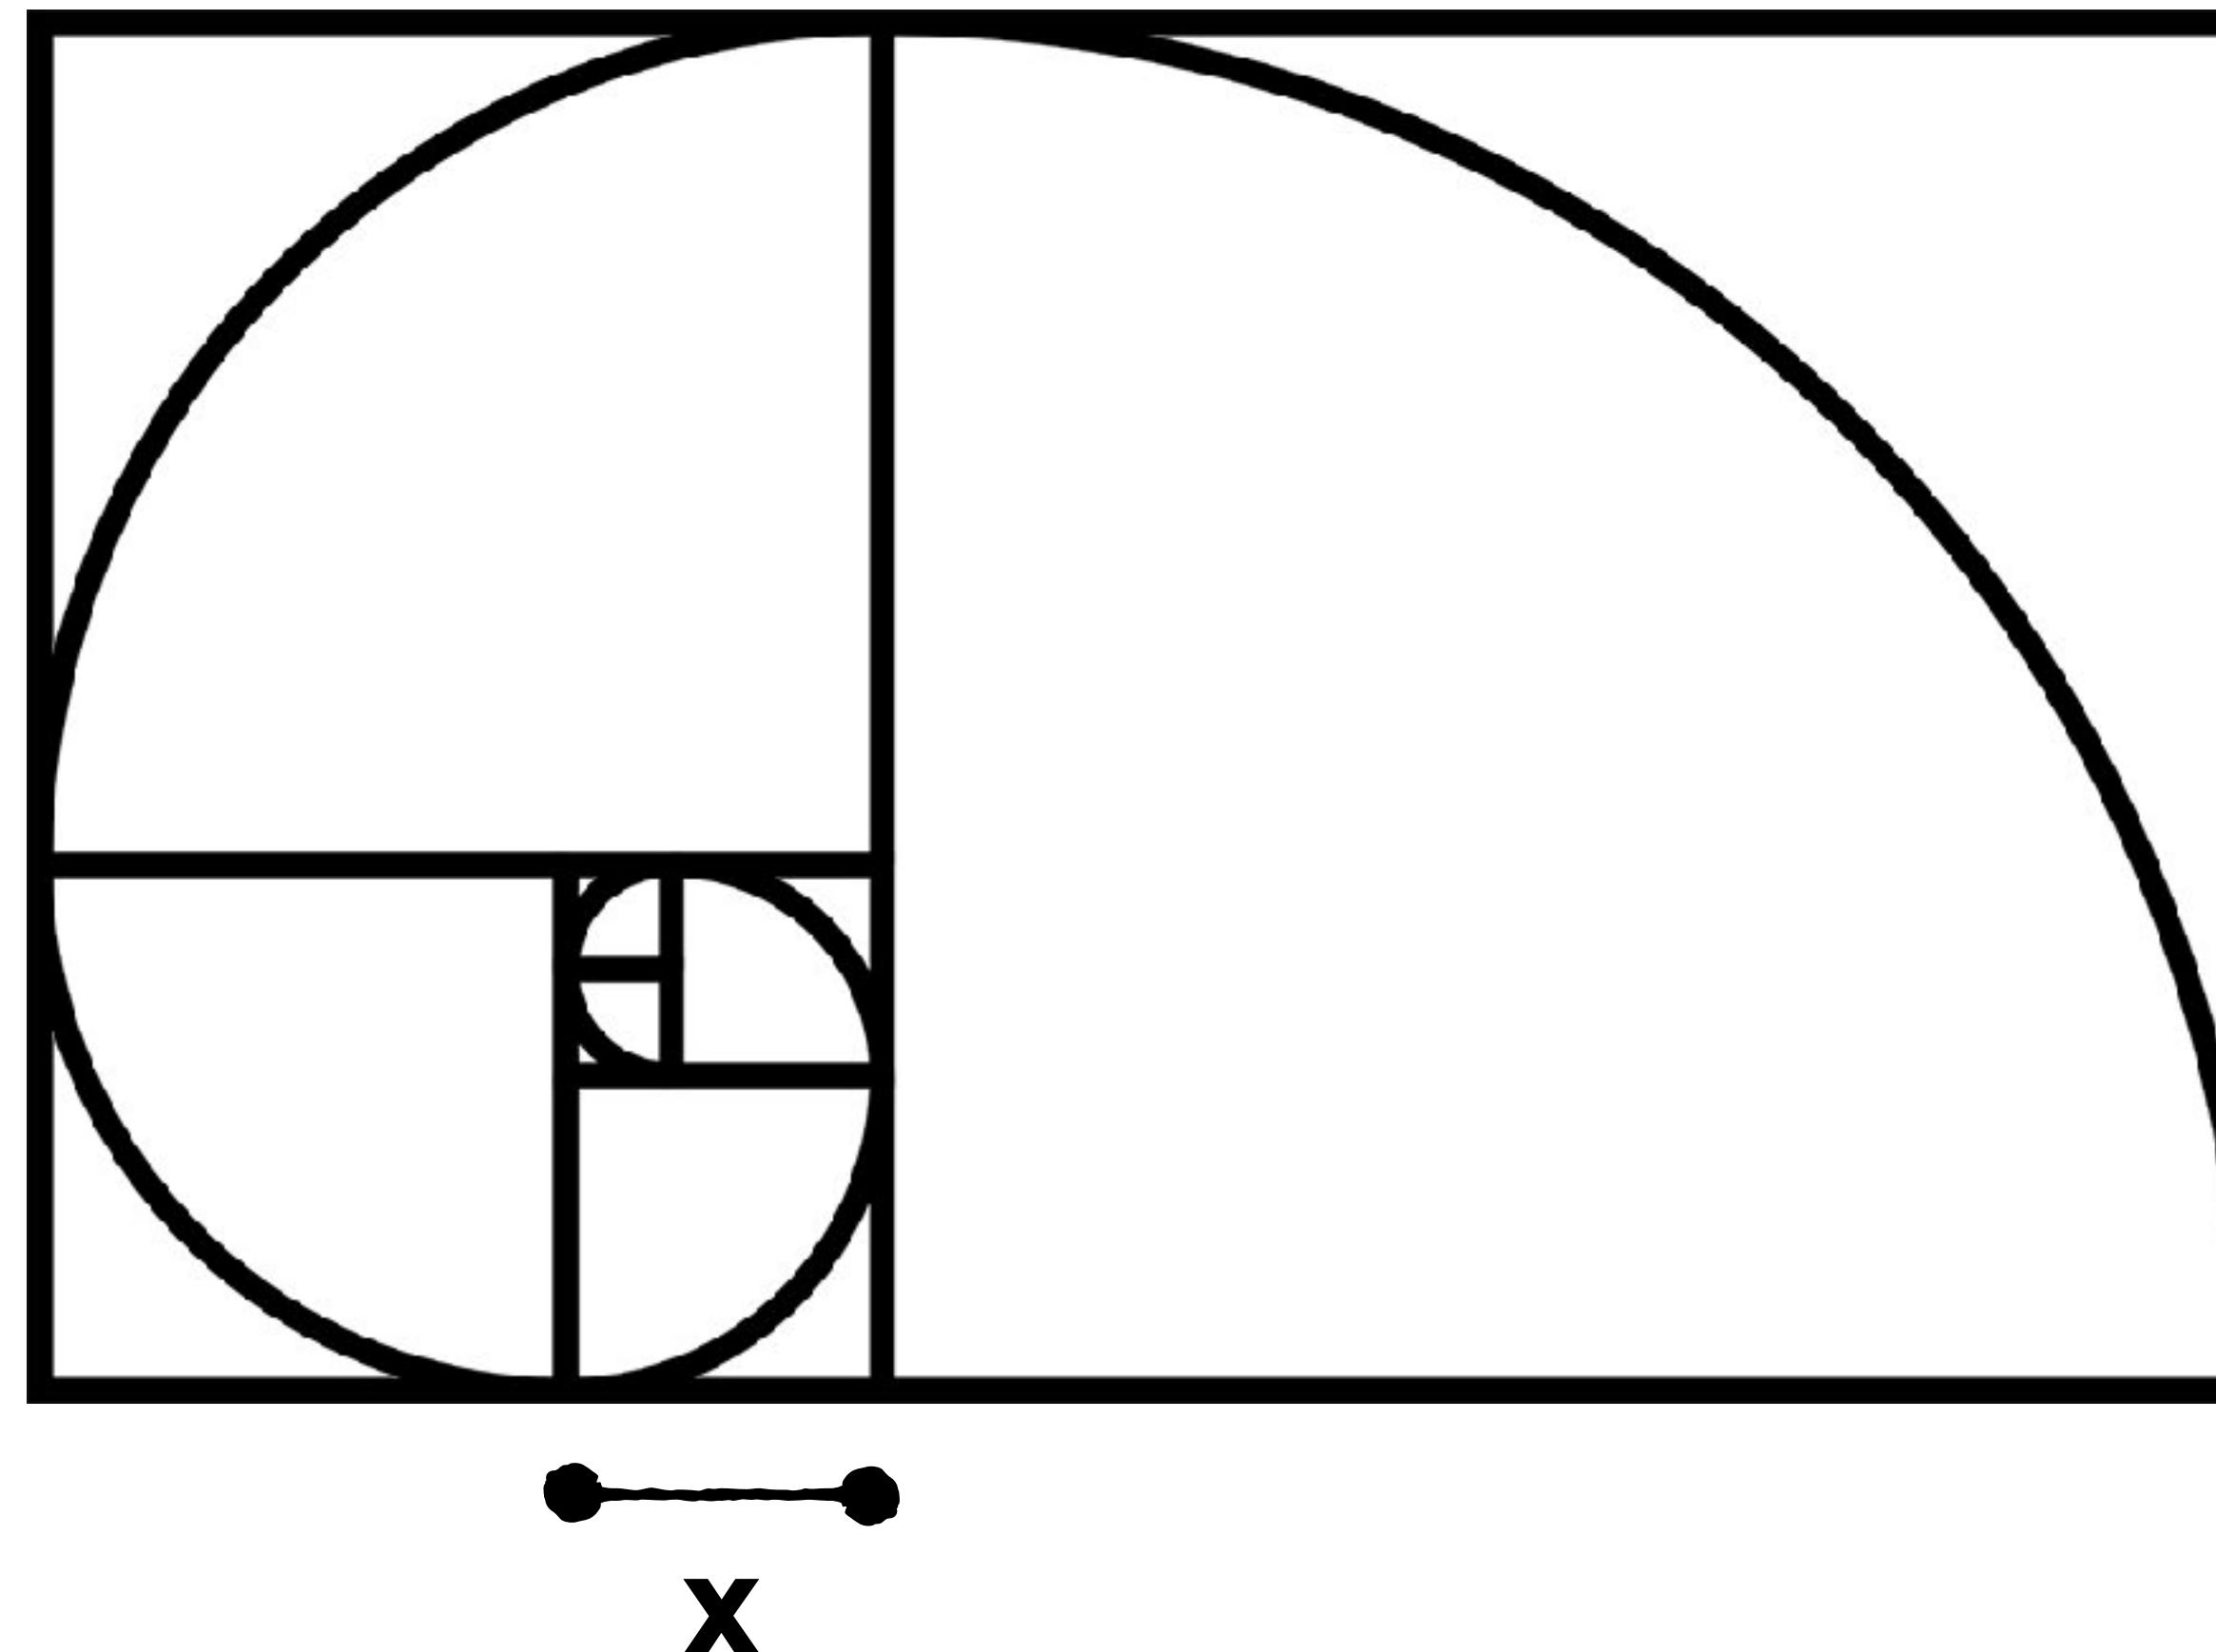
Ratio converges to the “Golden Ratio” $\phi = \frac{1 + \sqrt{5}}{2}$

Golden Spiral



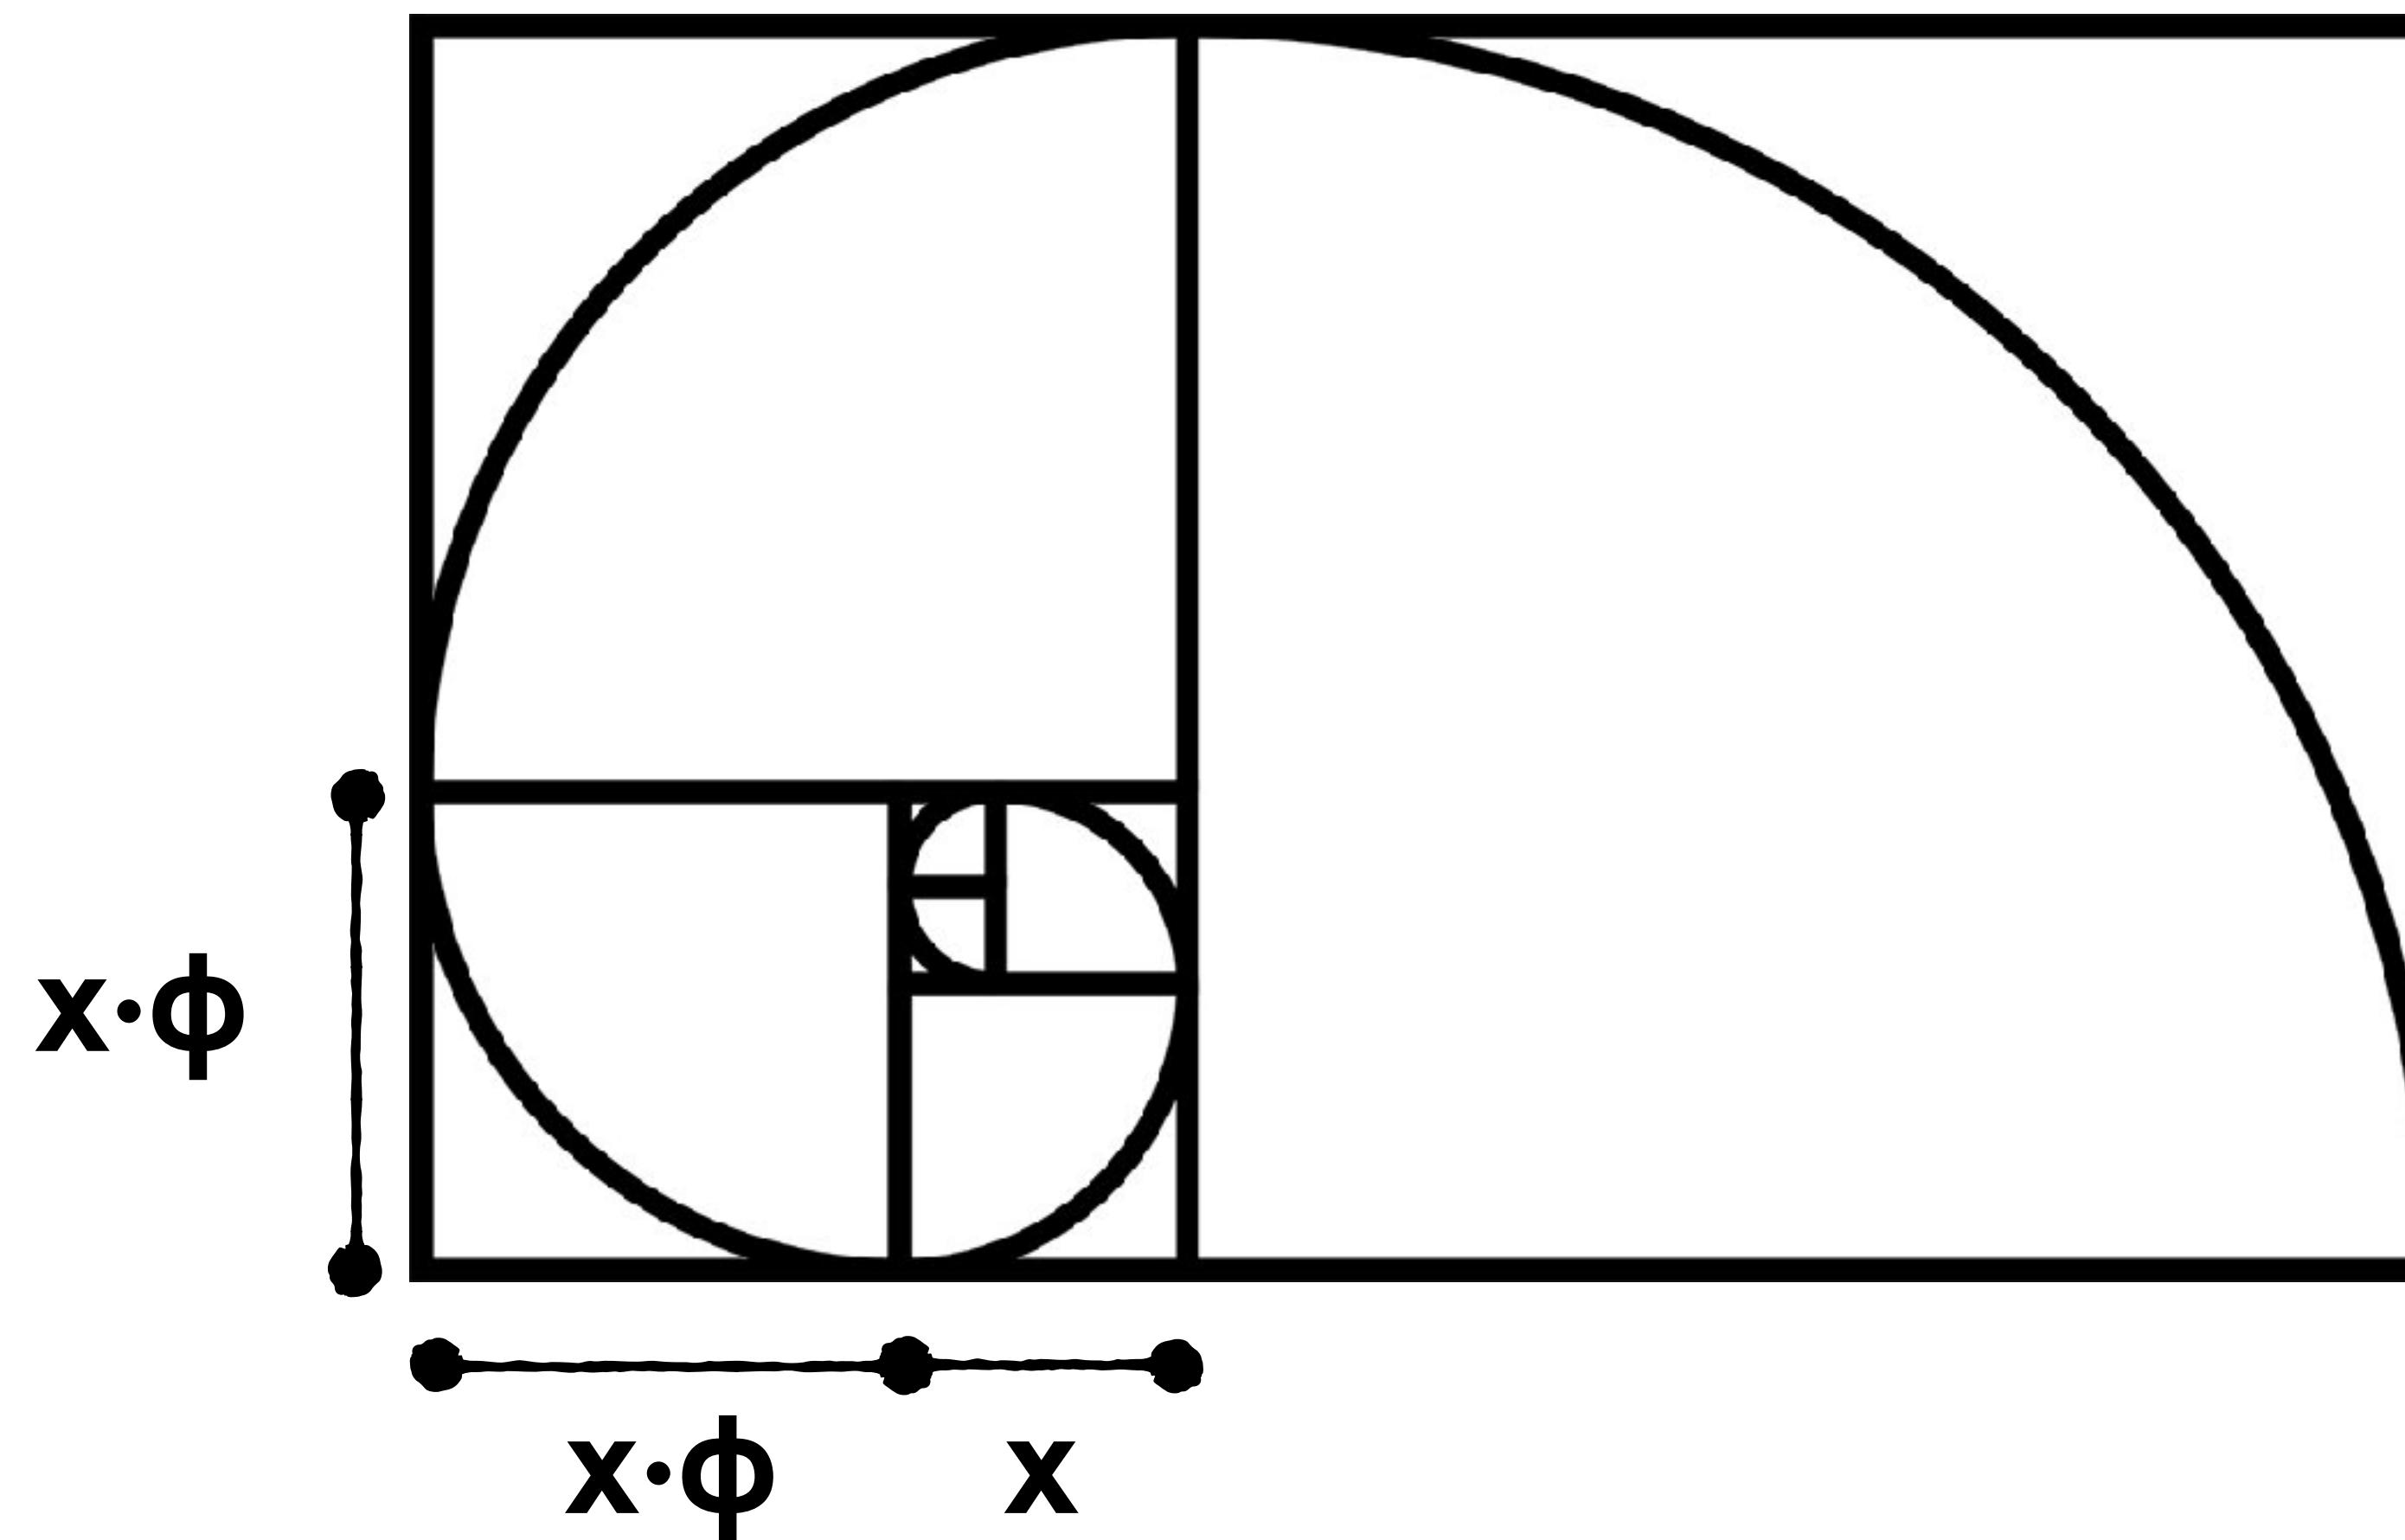
Ratio converges to the “Golden Ratio” $\phi = \frac{1 + \sqrt{5}}{2}$

Golden Spiral

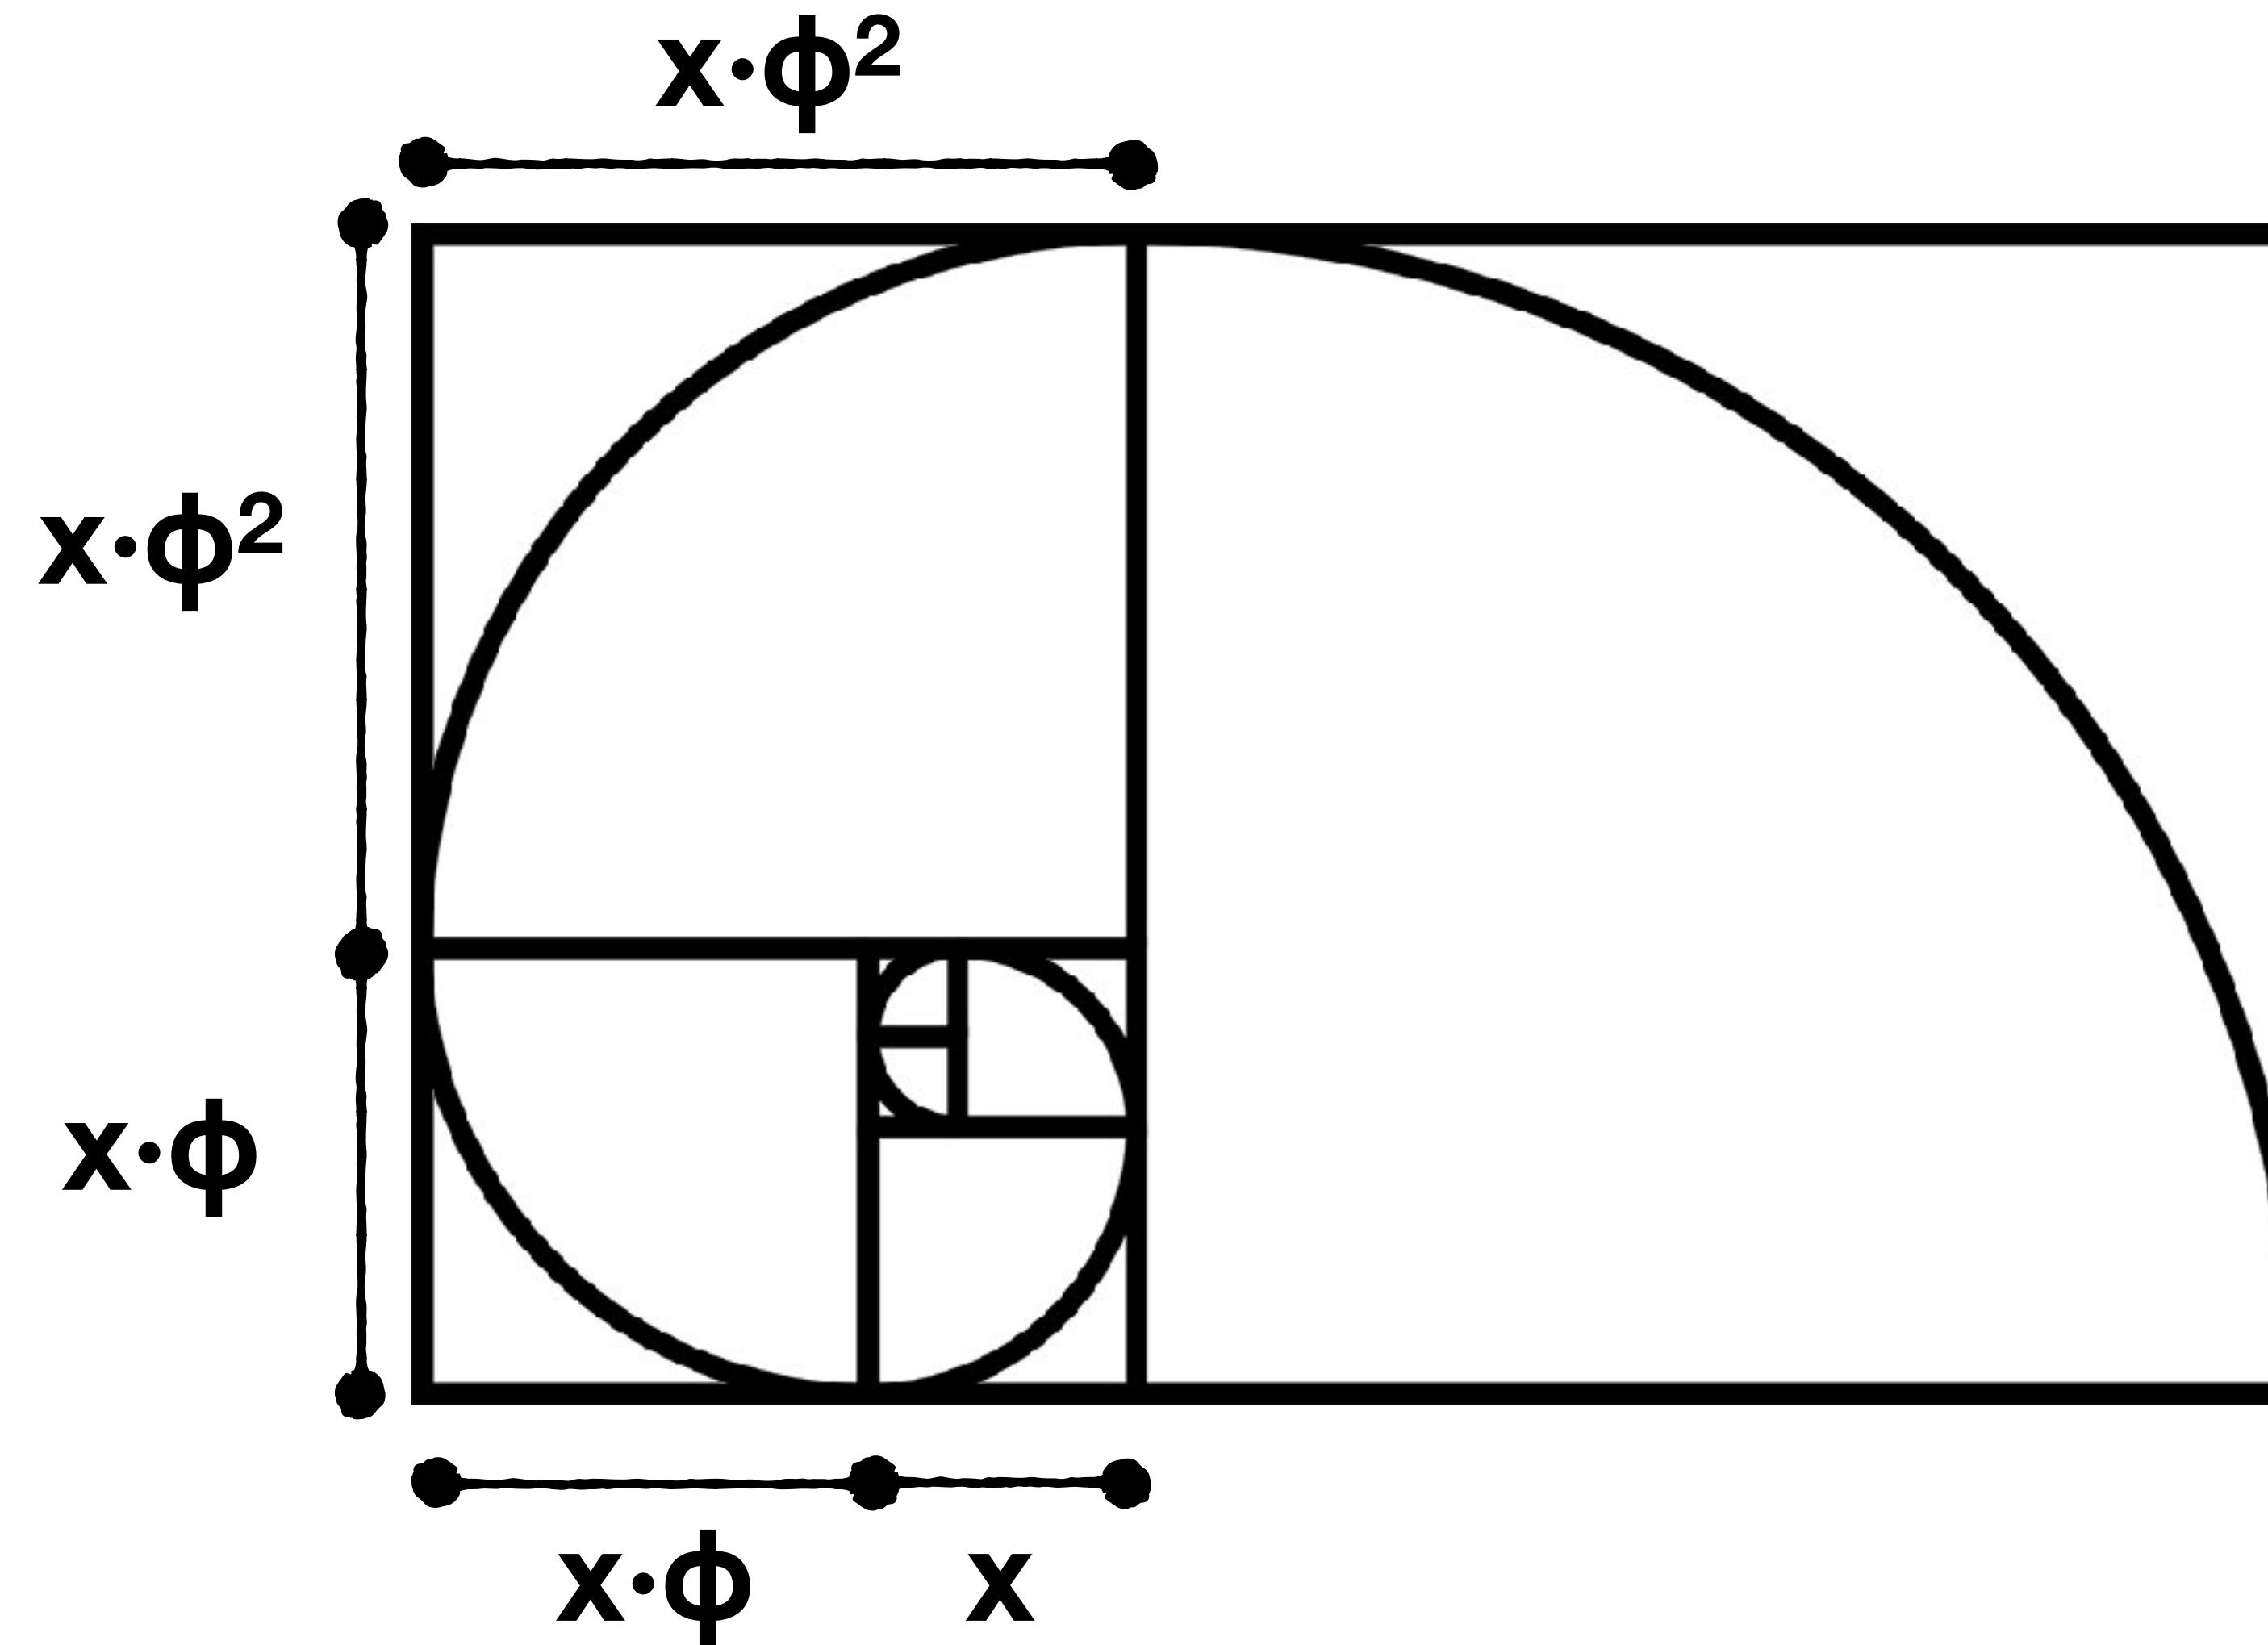


Ratio converges to the “Golden Ratio” $\phi = \frac{1 + \sqrt{5}}{2}$

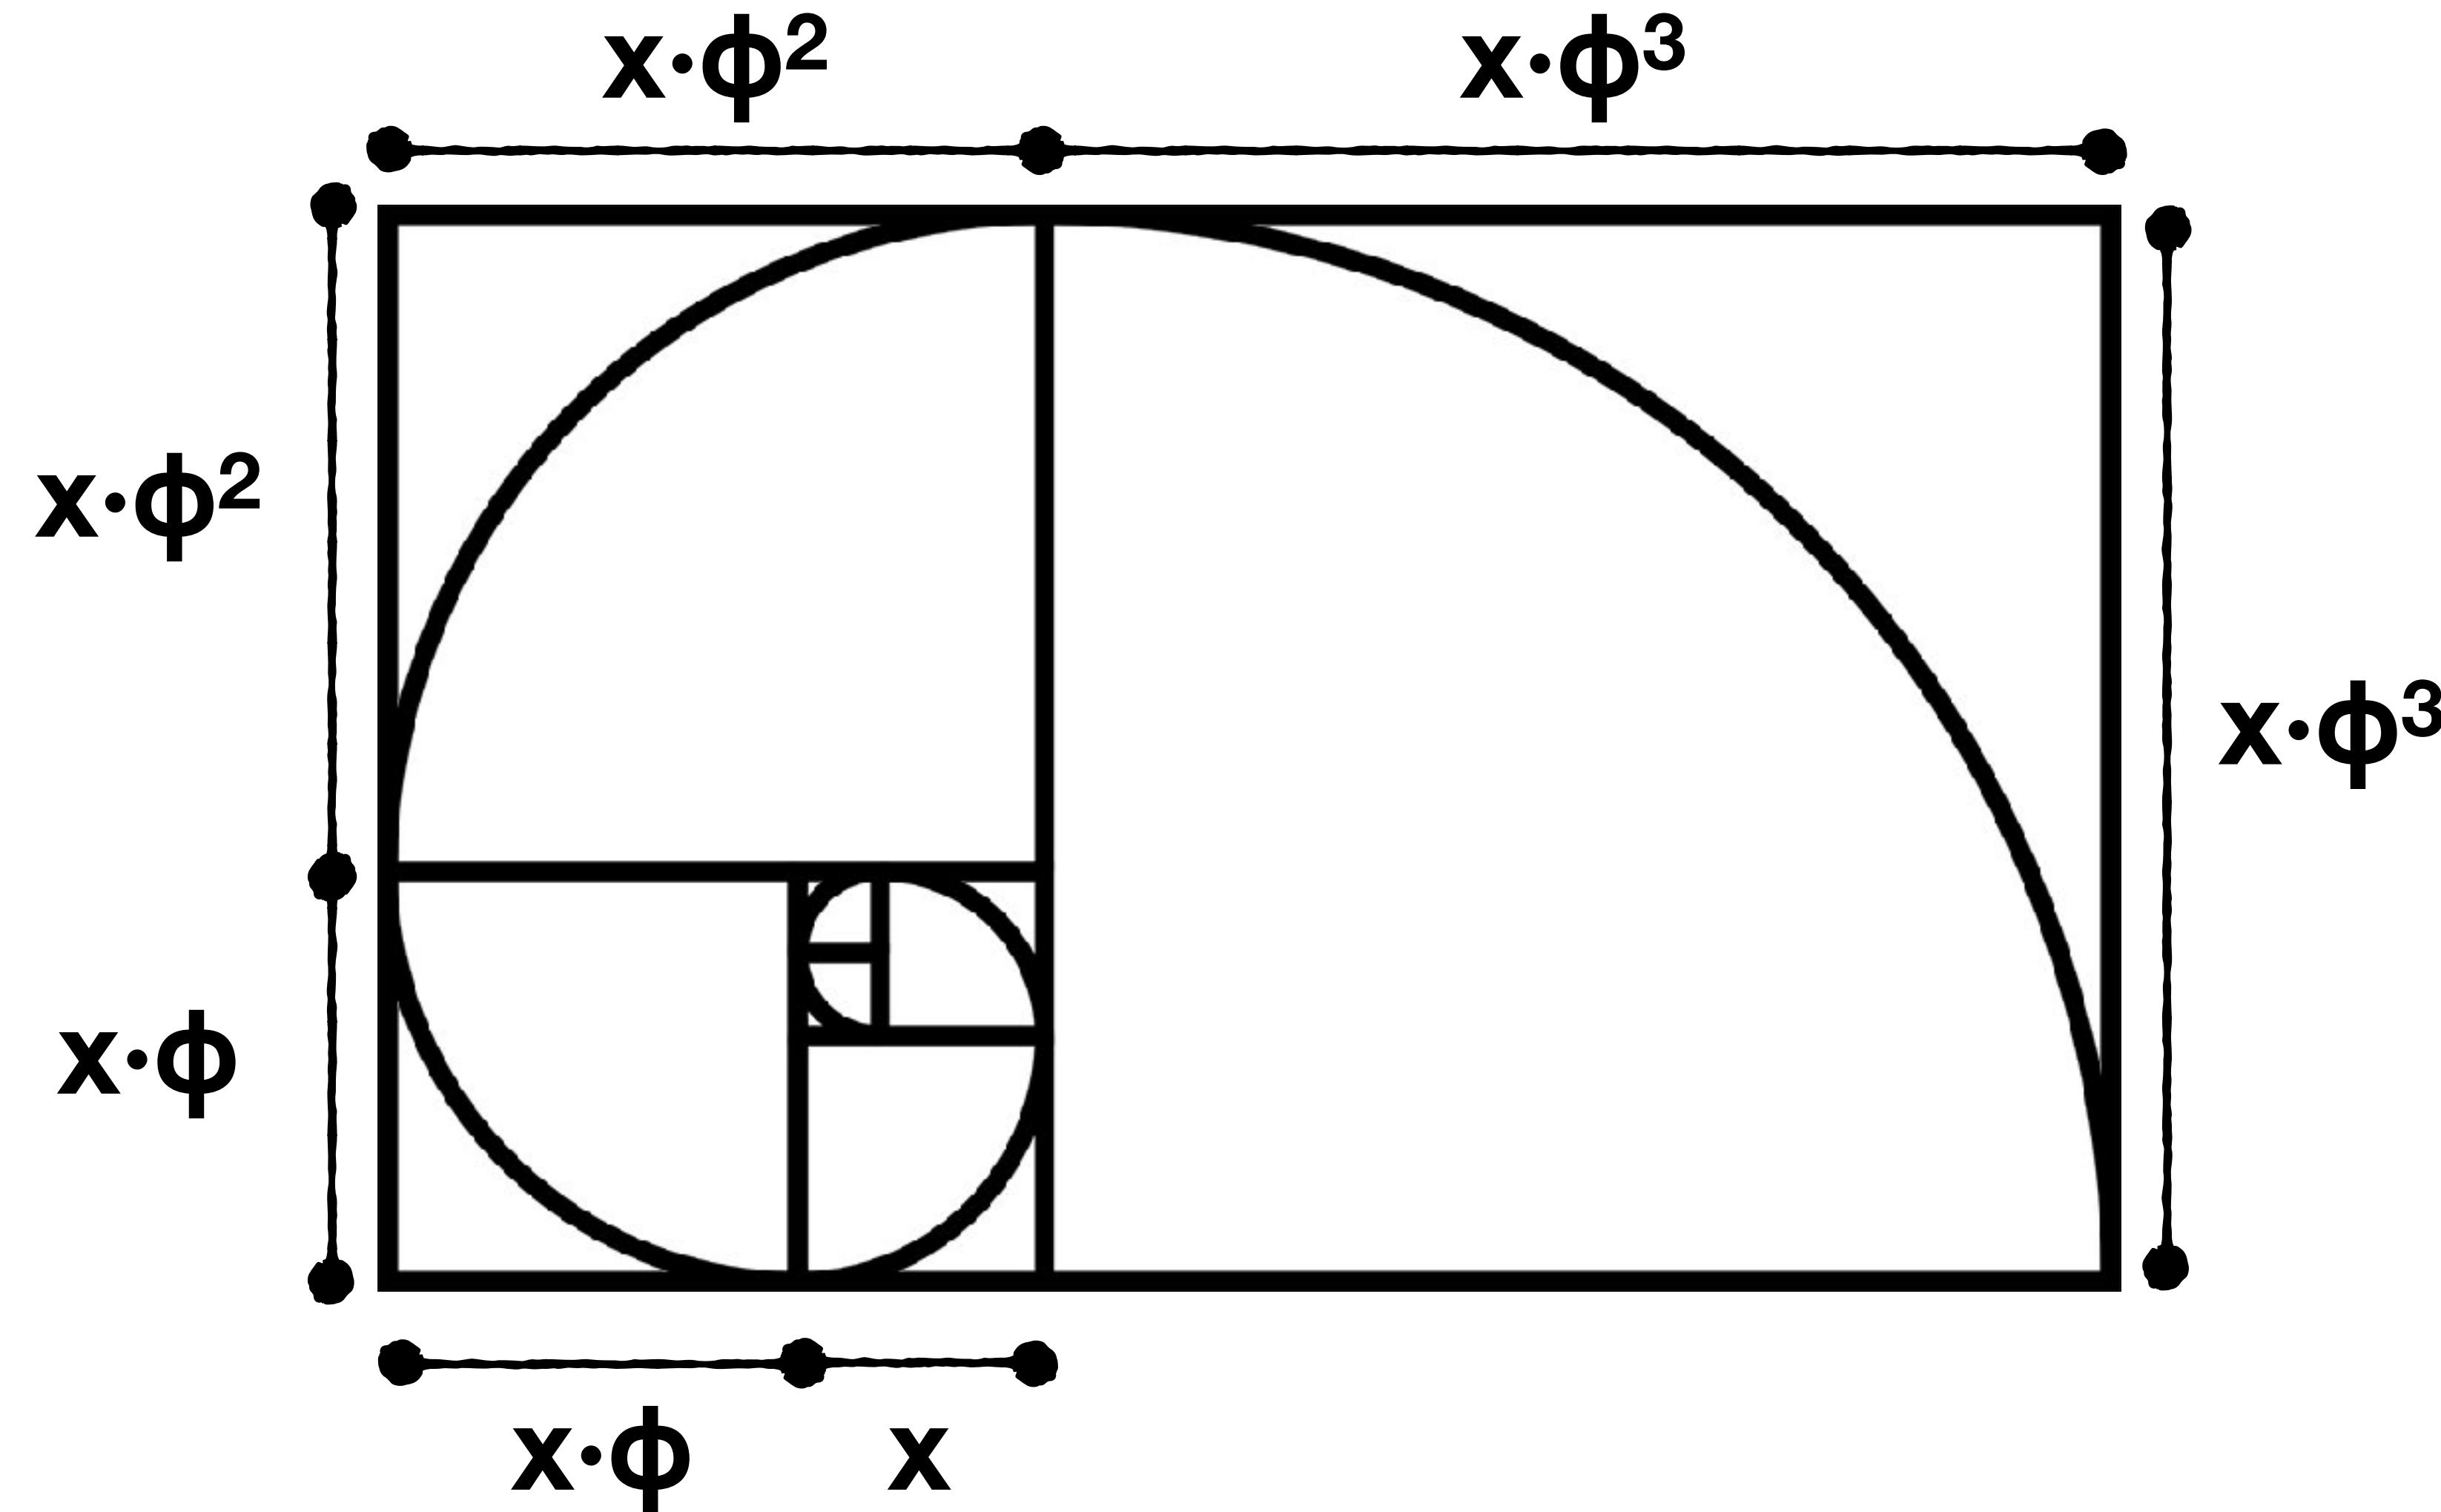
Golden Spiral



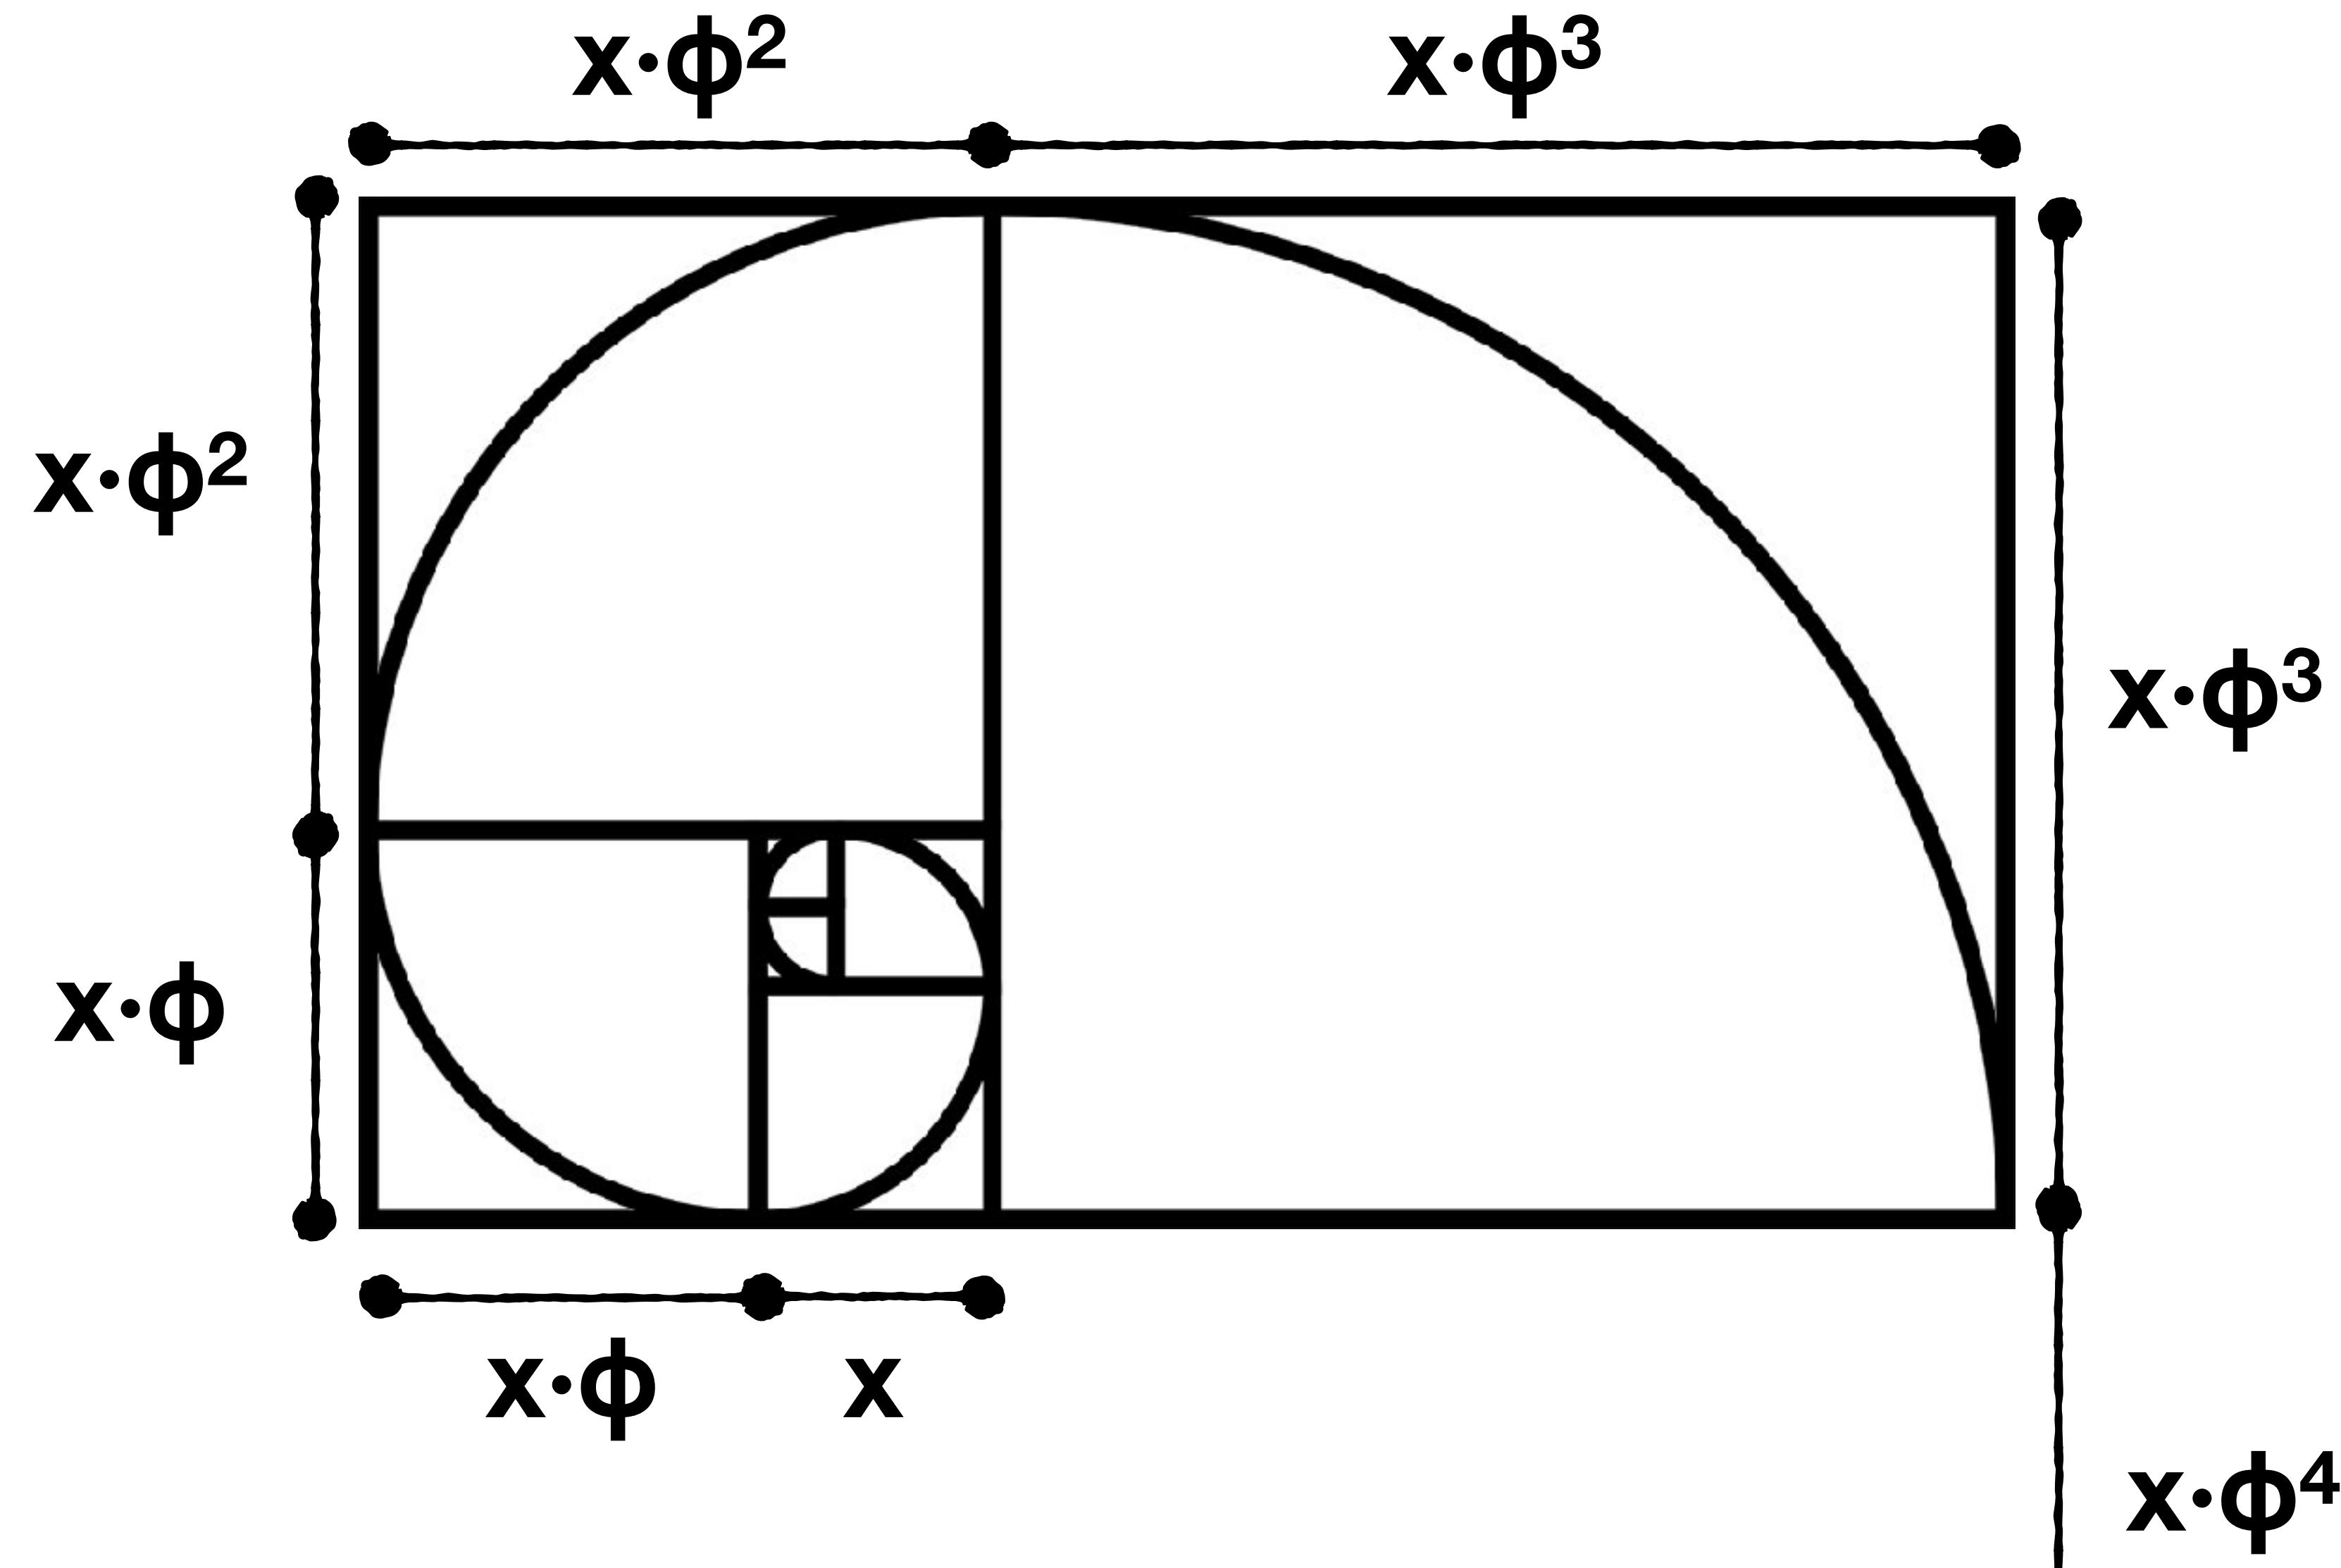
Golden Spiral



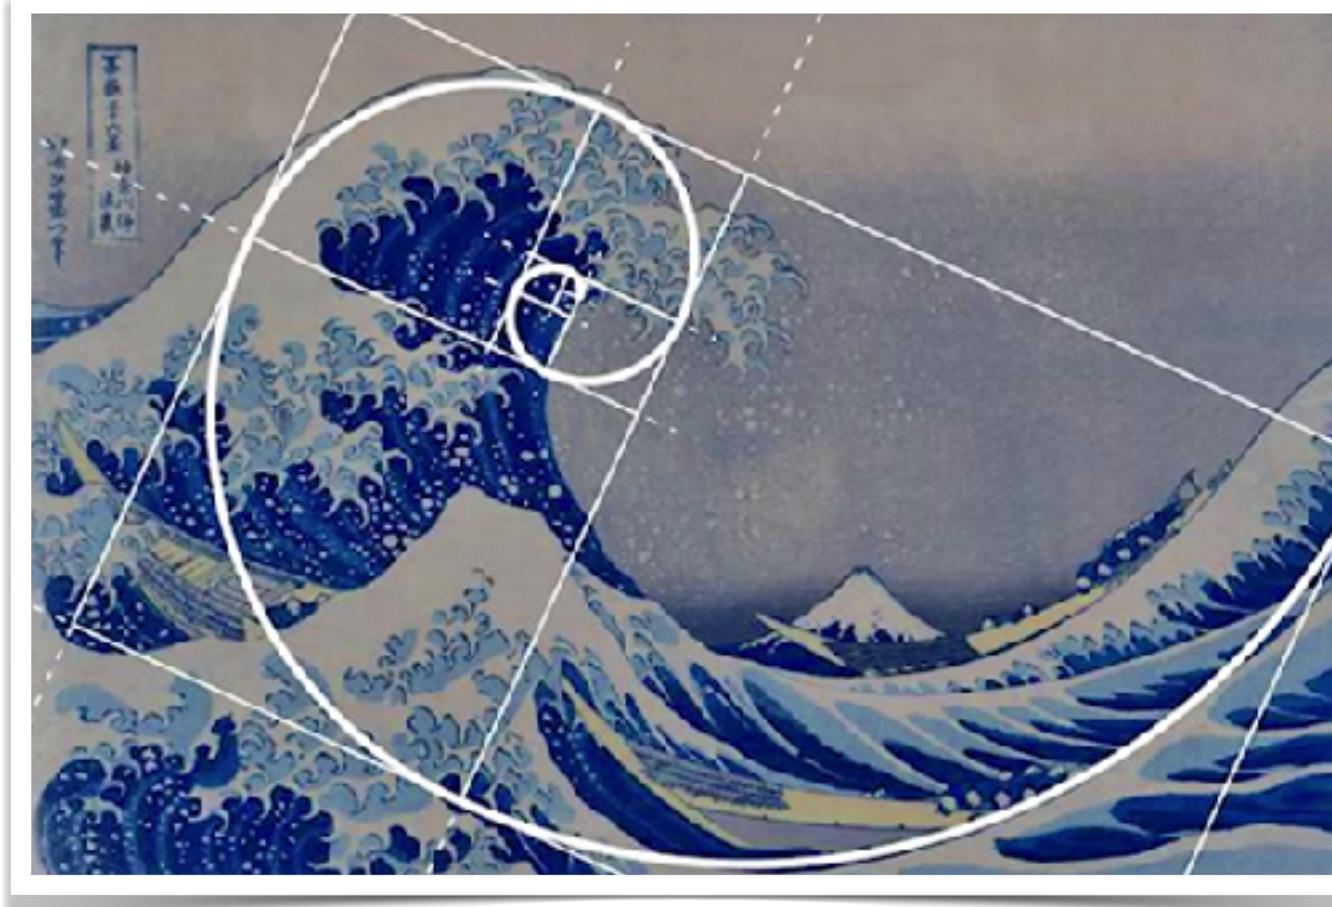
Golden Spiral



Golden Spiral



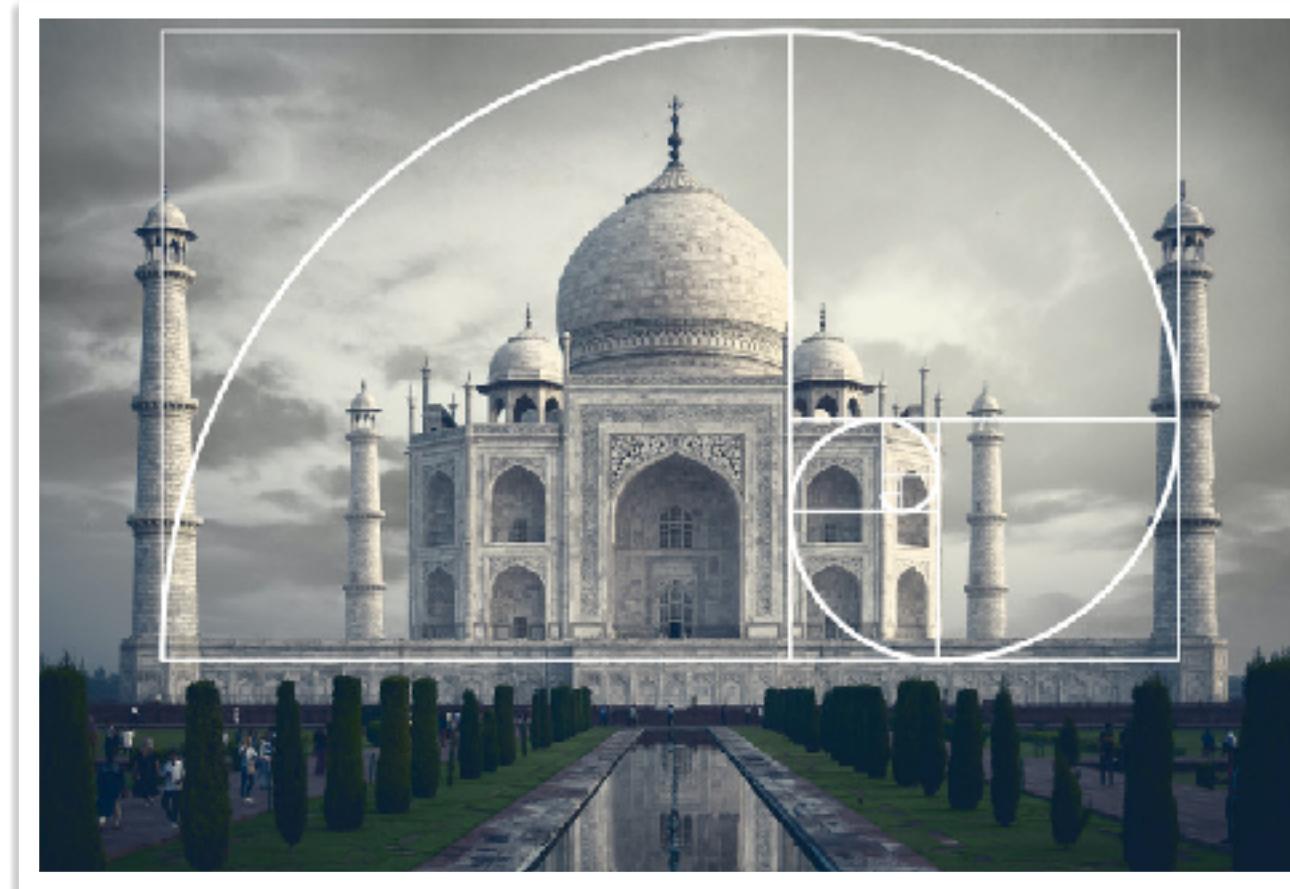
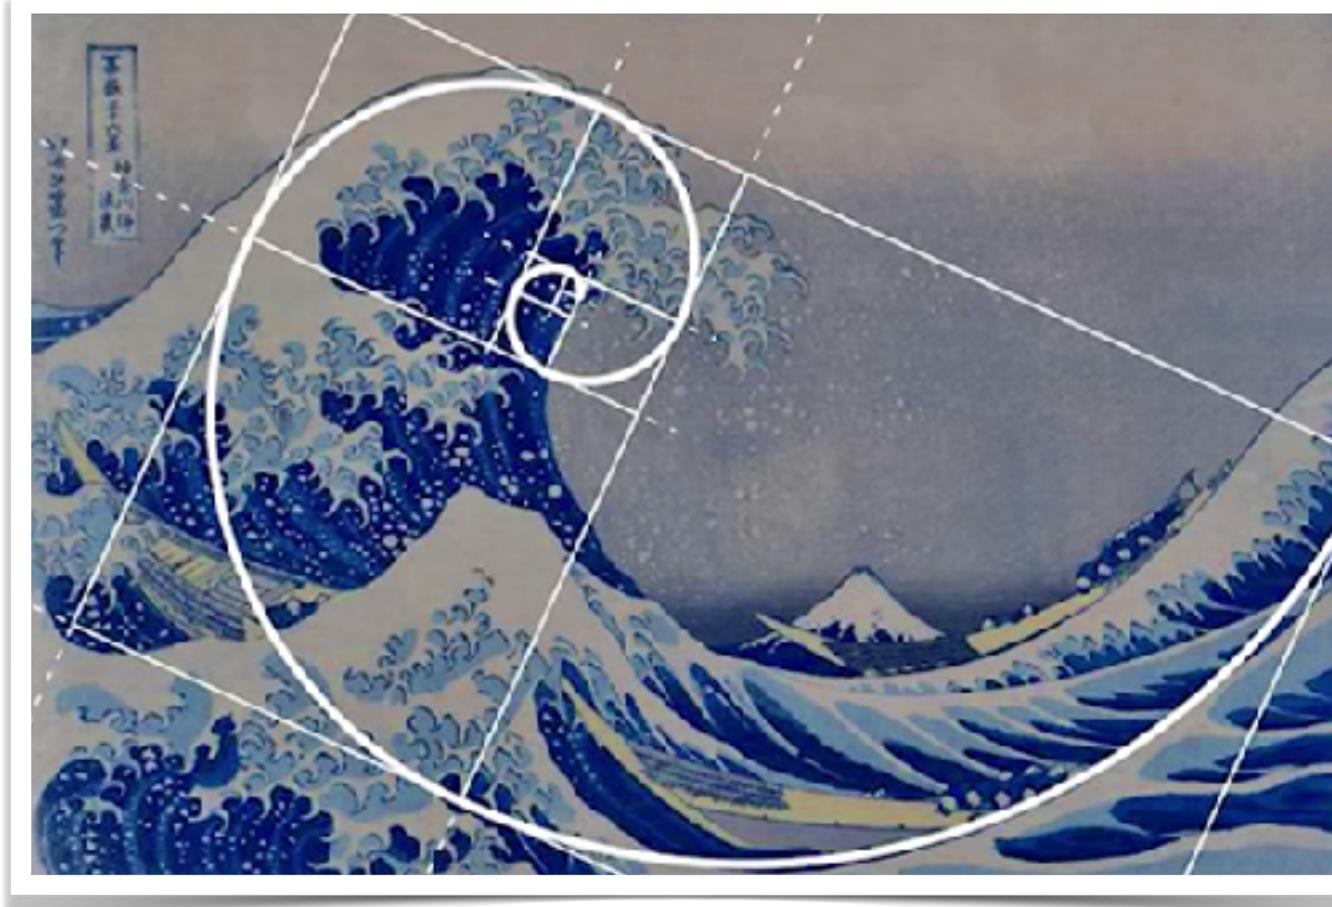
Golden Spiral



Art

**The Great Wave
off Kanagawa**

Golden Spiral



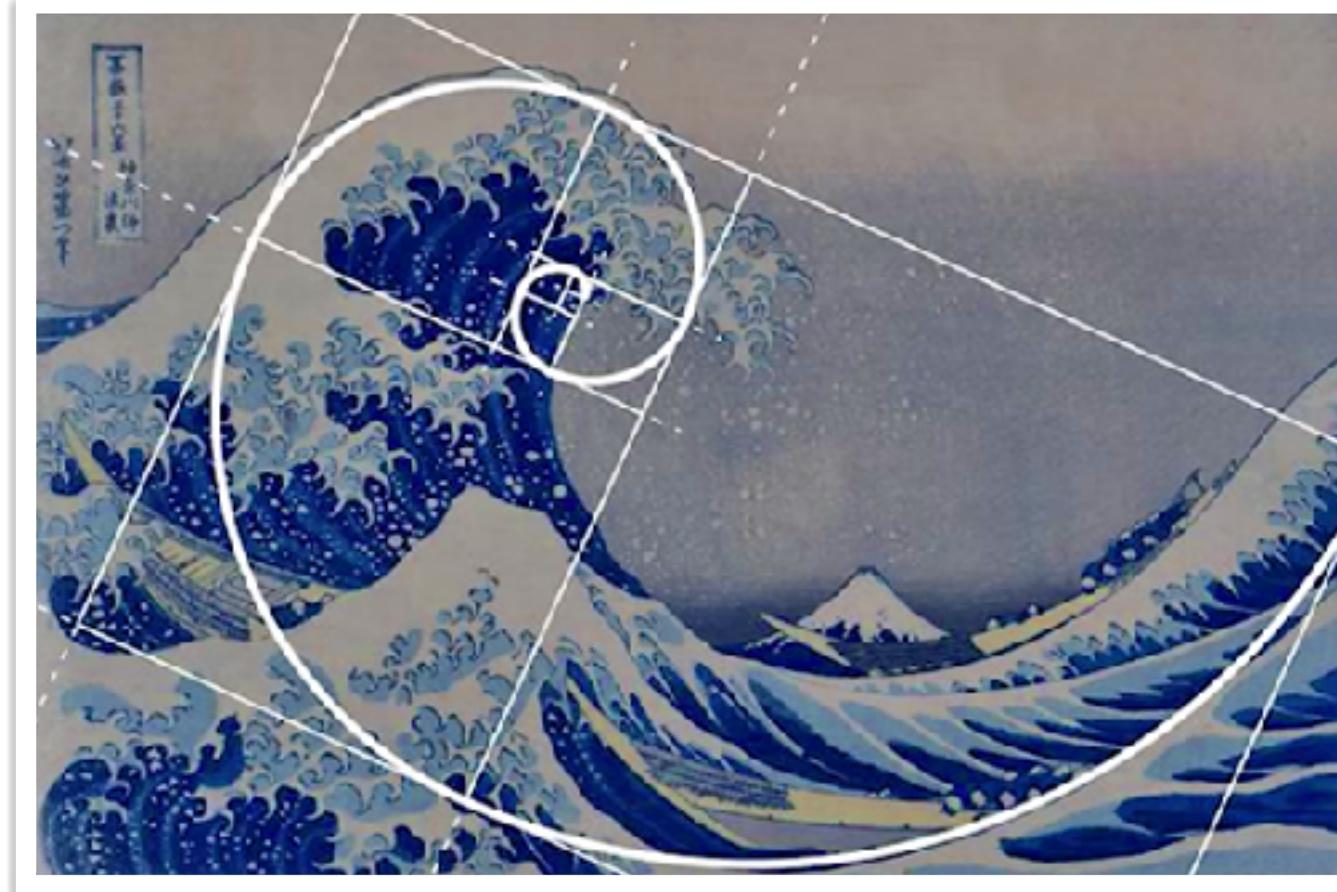
Art

The Great Wave
off Kanagawa

Architecture

Taj Mahal

Golden Spiral



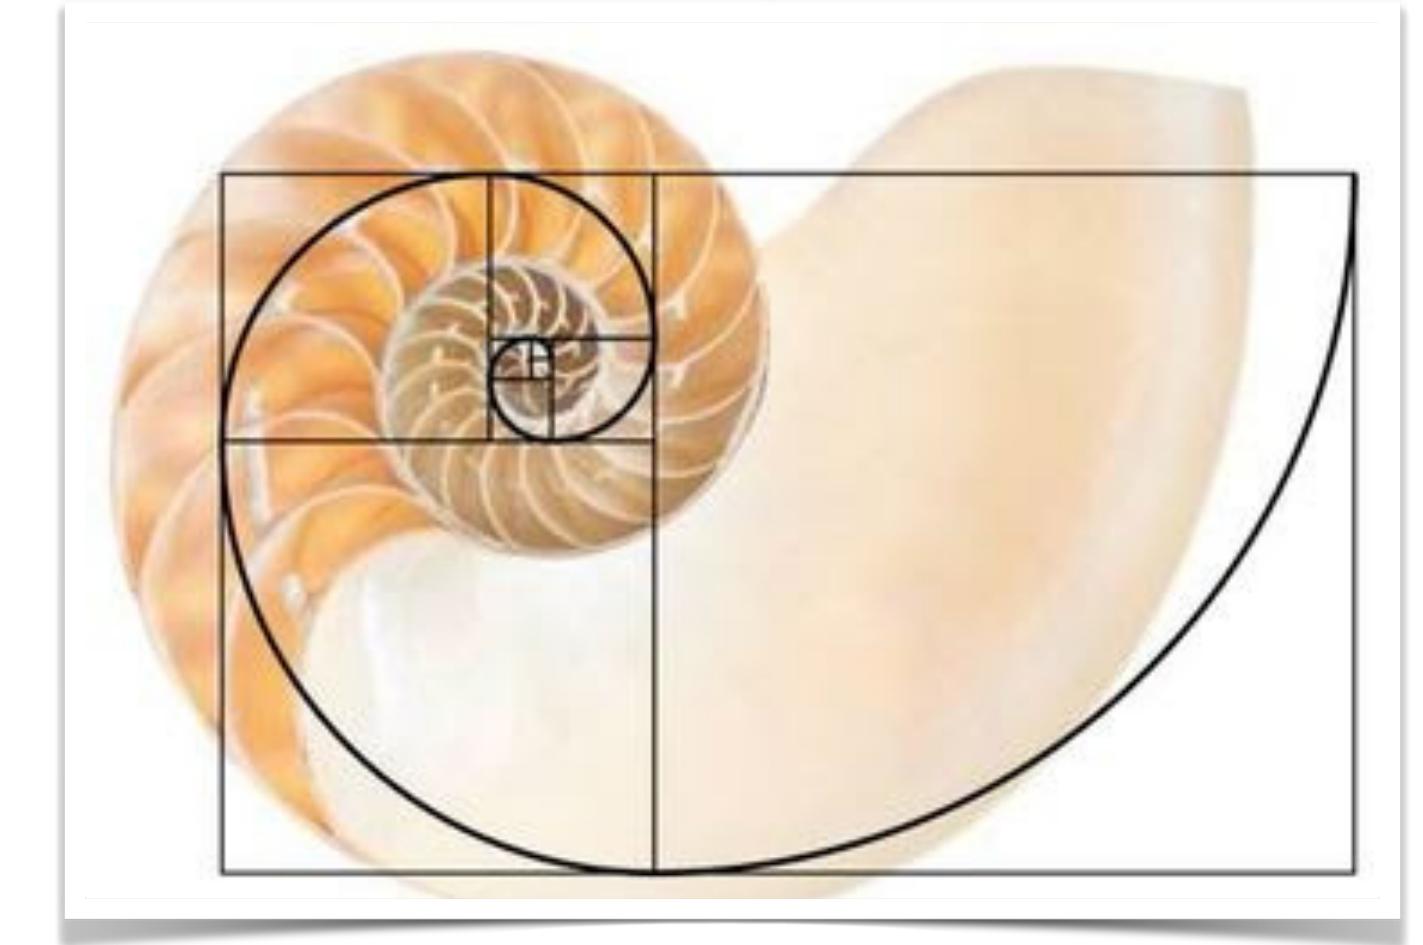
Art

The Great Wave
off Kanagawa



Architecture

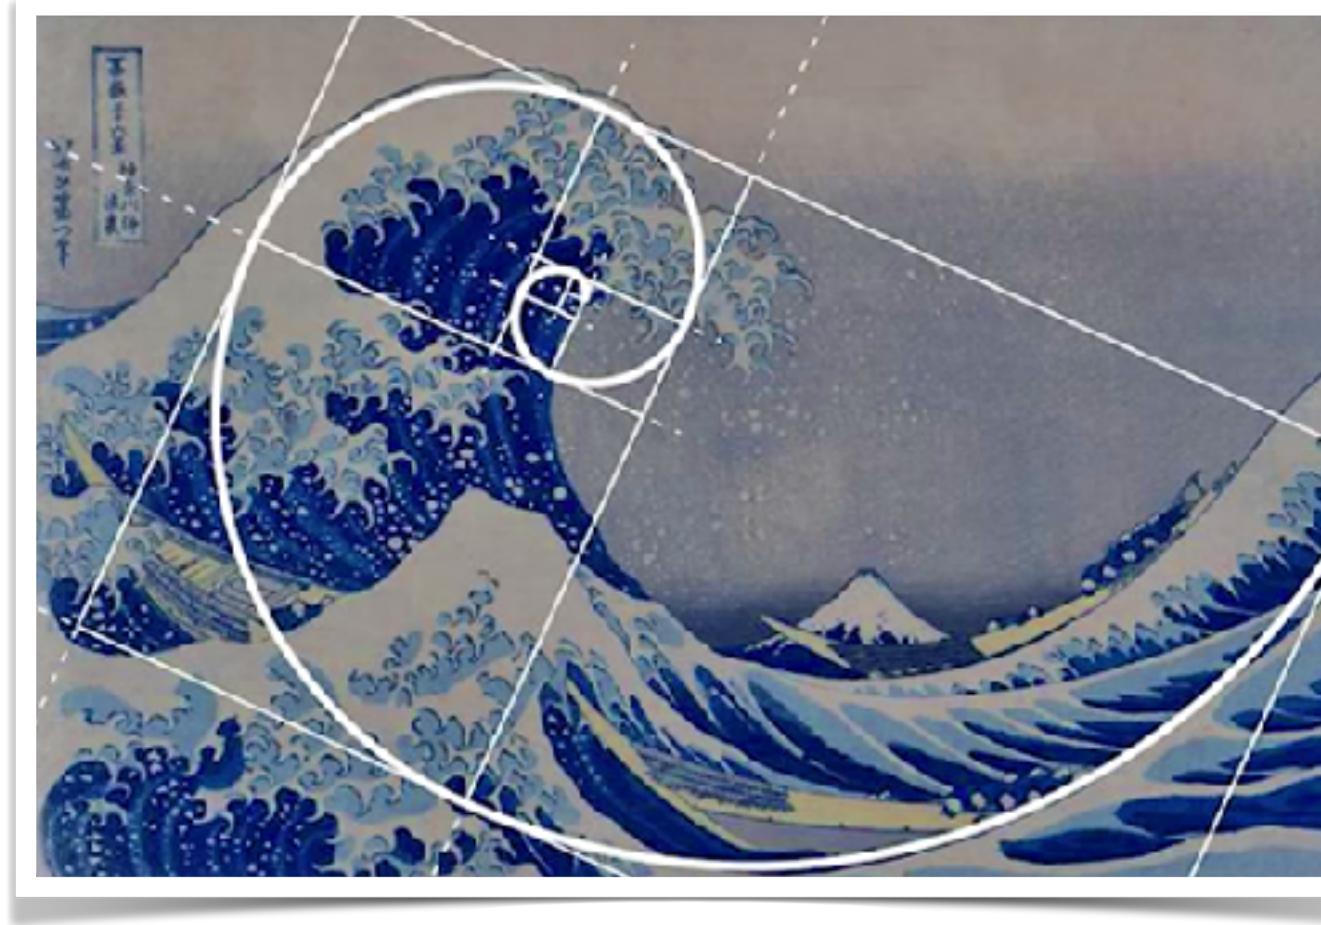
Taj Mahal



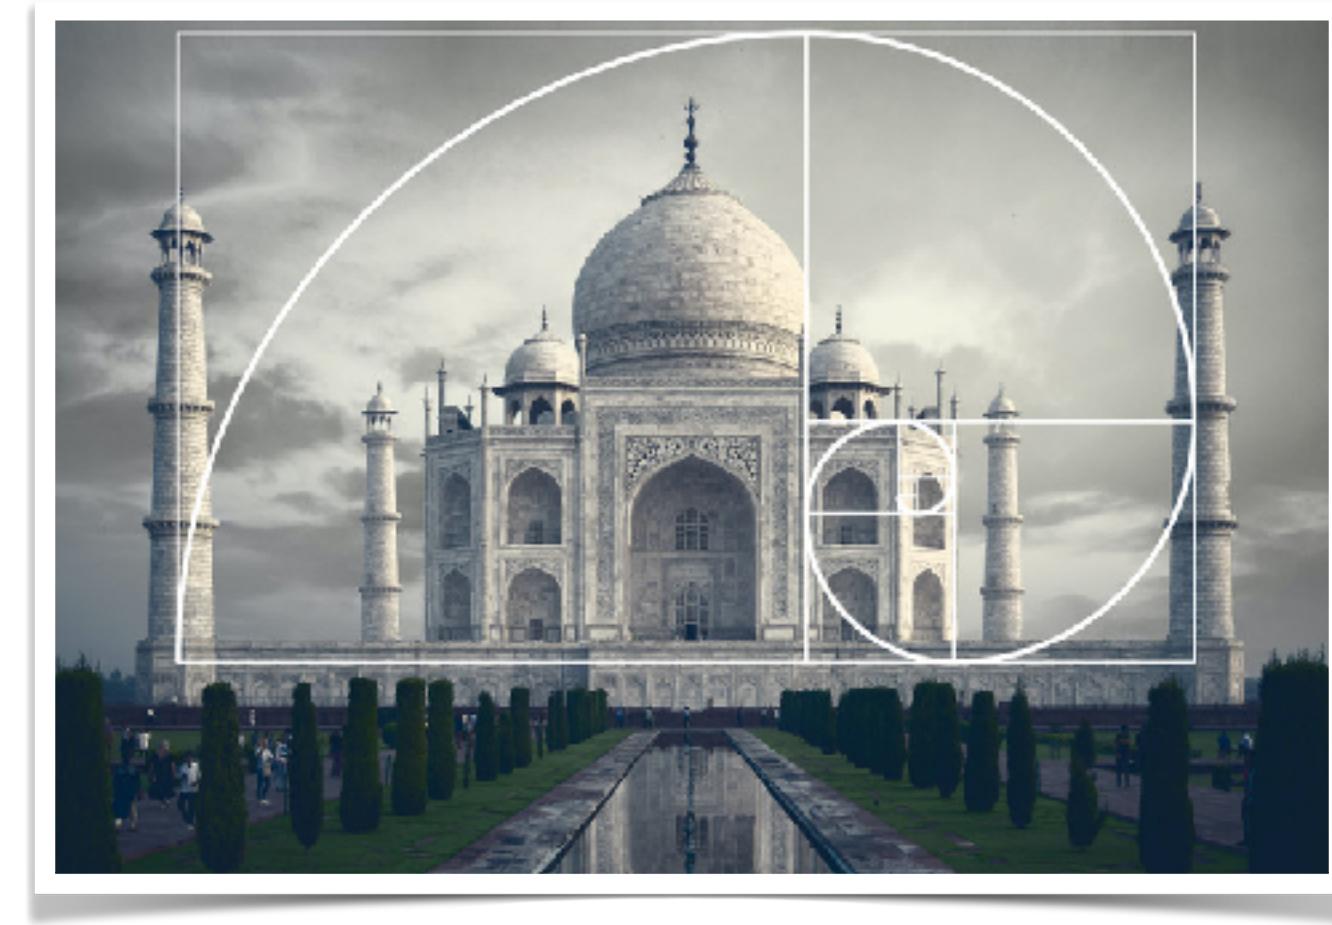
Nature

Nautilus Shell

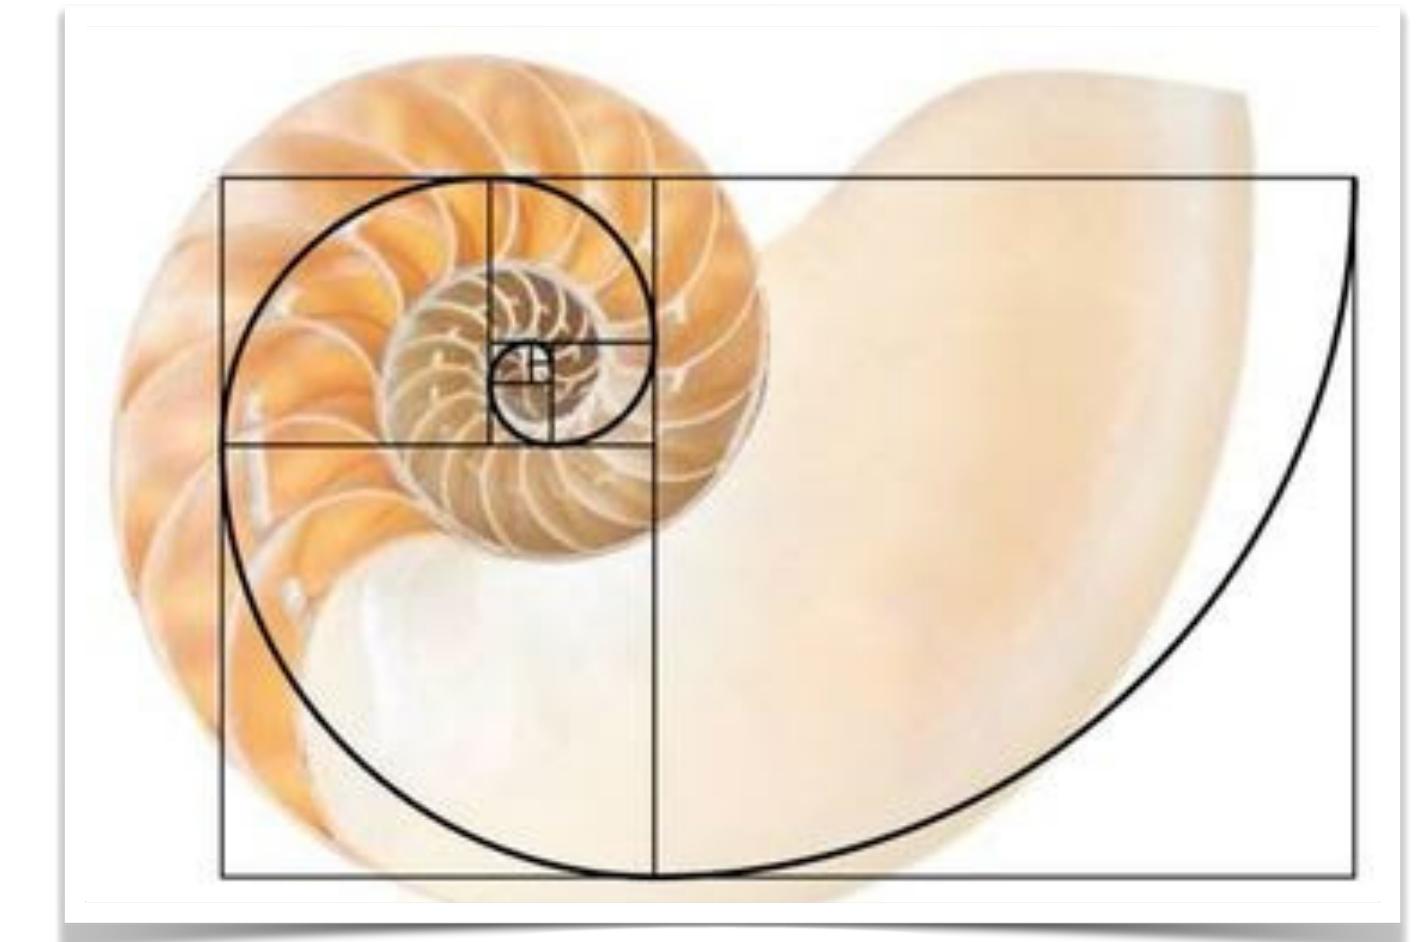
Golden Spiral



Art



Architecture



Nature

And now also in computer science :)

Weird Properties

$$\phi = \phi^2 - 1 = \frac{1}{\phi - 1}$$

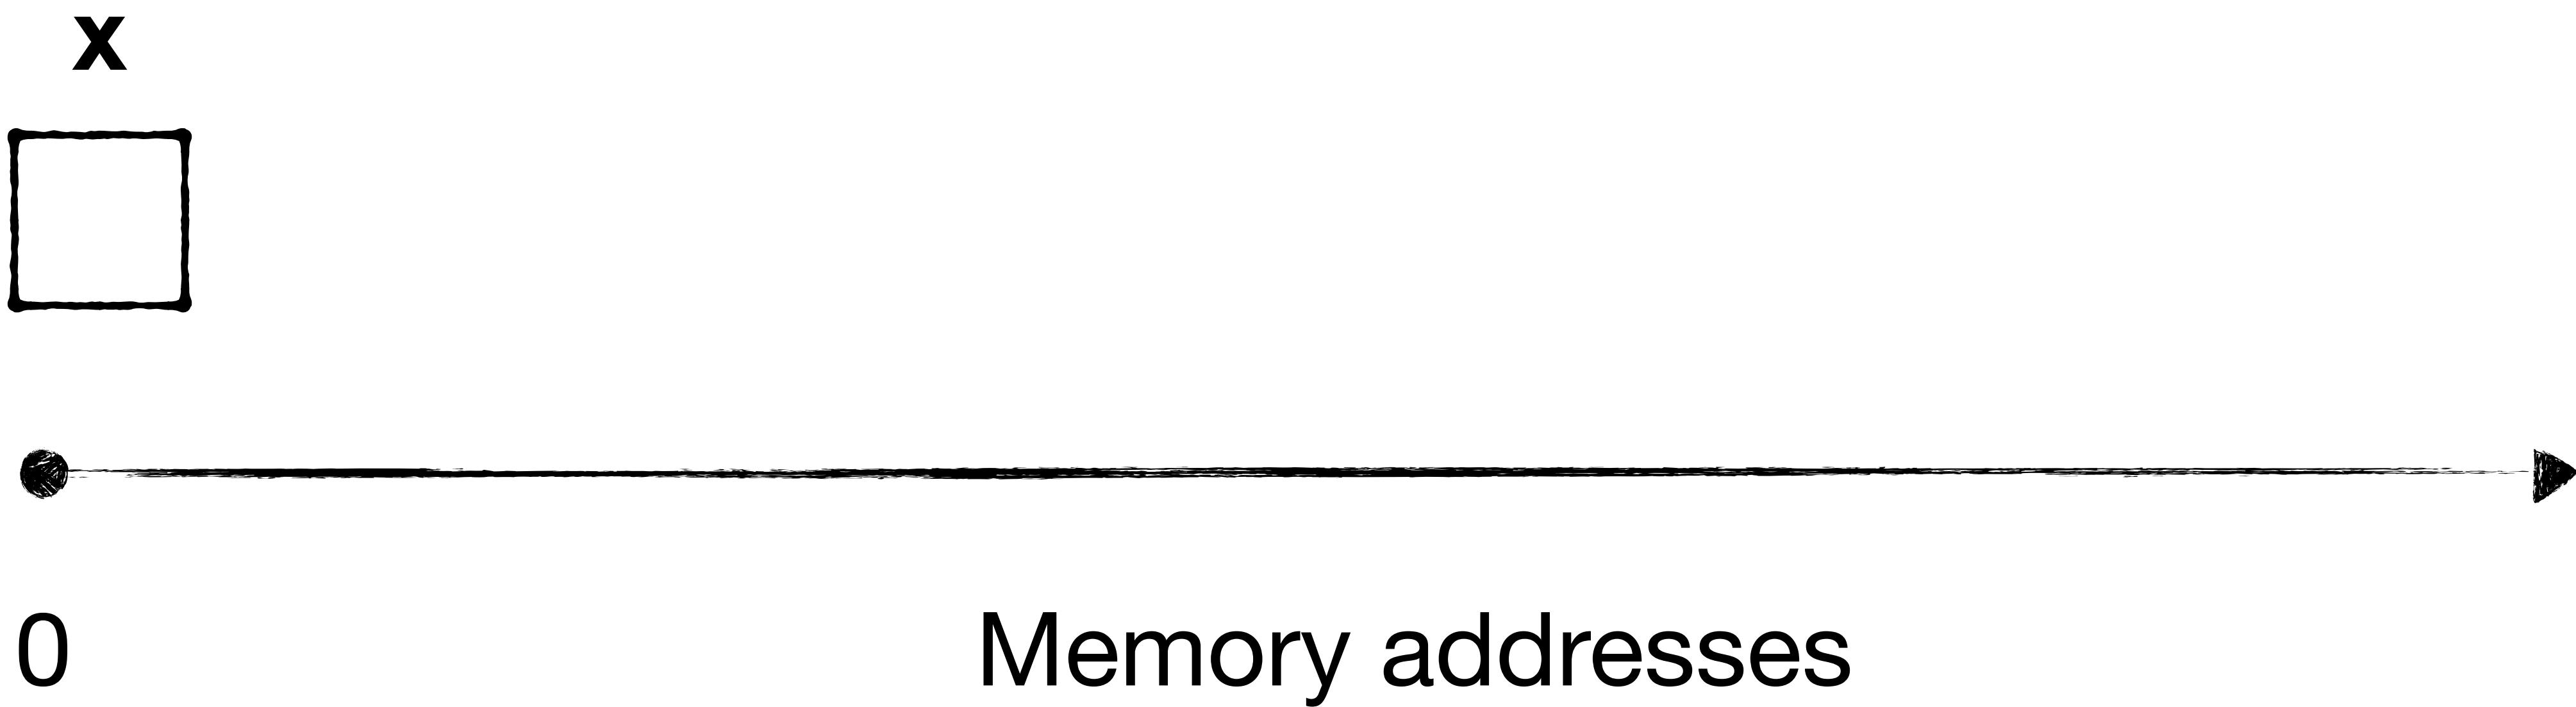
Expand Array by Golden Ratio ($G = \phi = 1.61\dots$)

Expand Array by Golden Ratio ($G = \phi = 1.61\dots$)

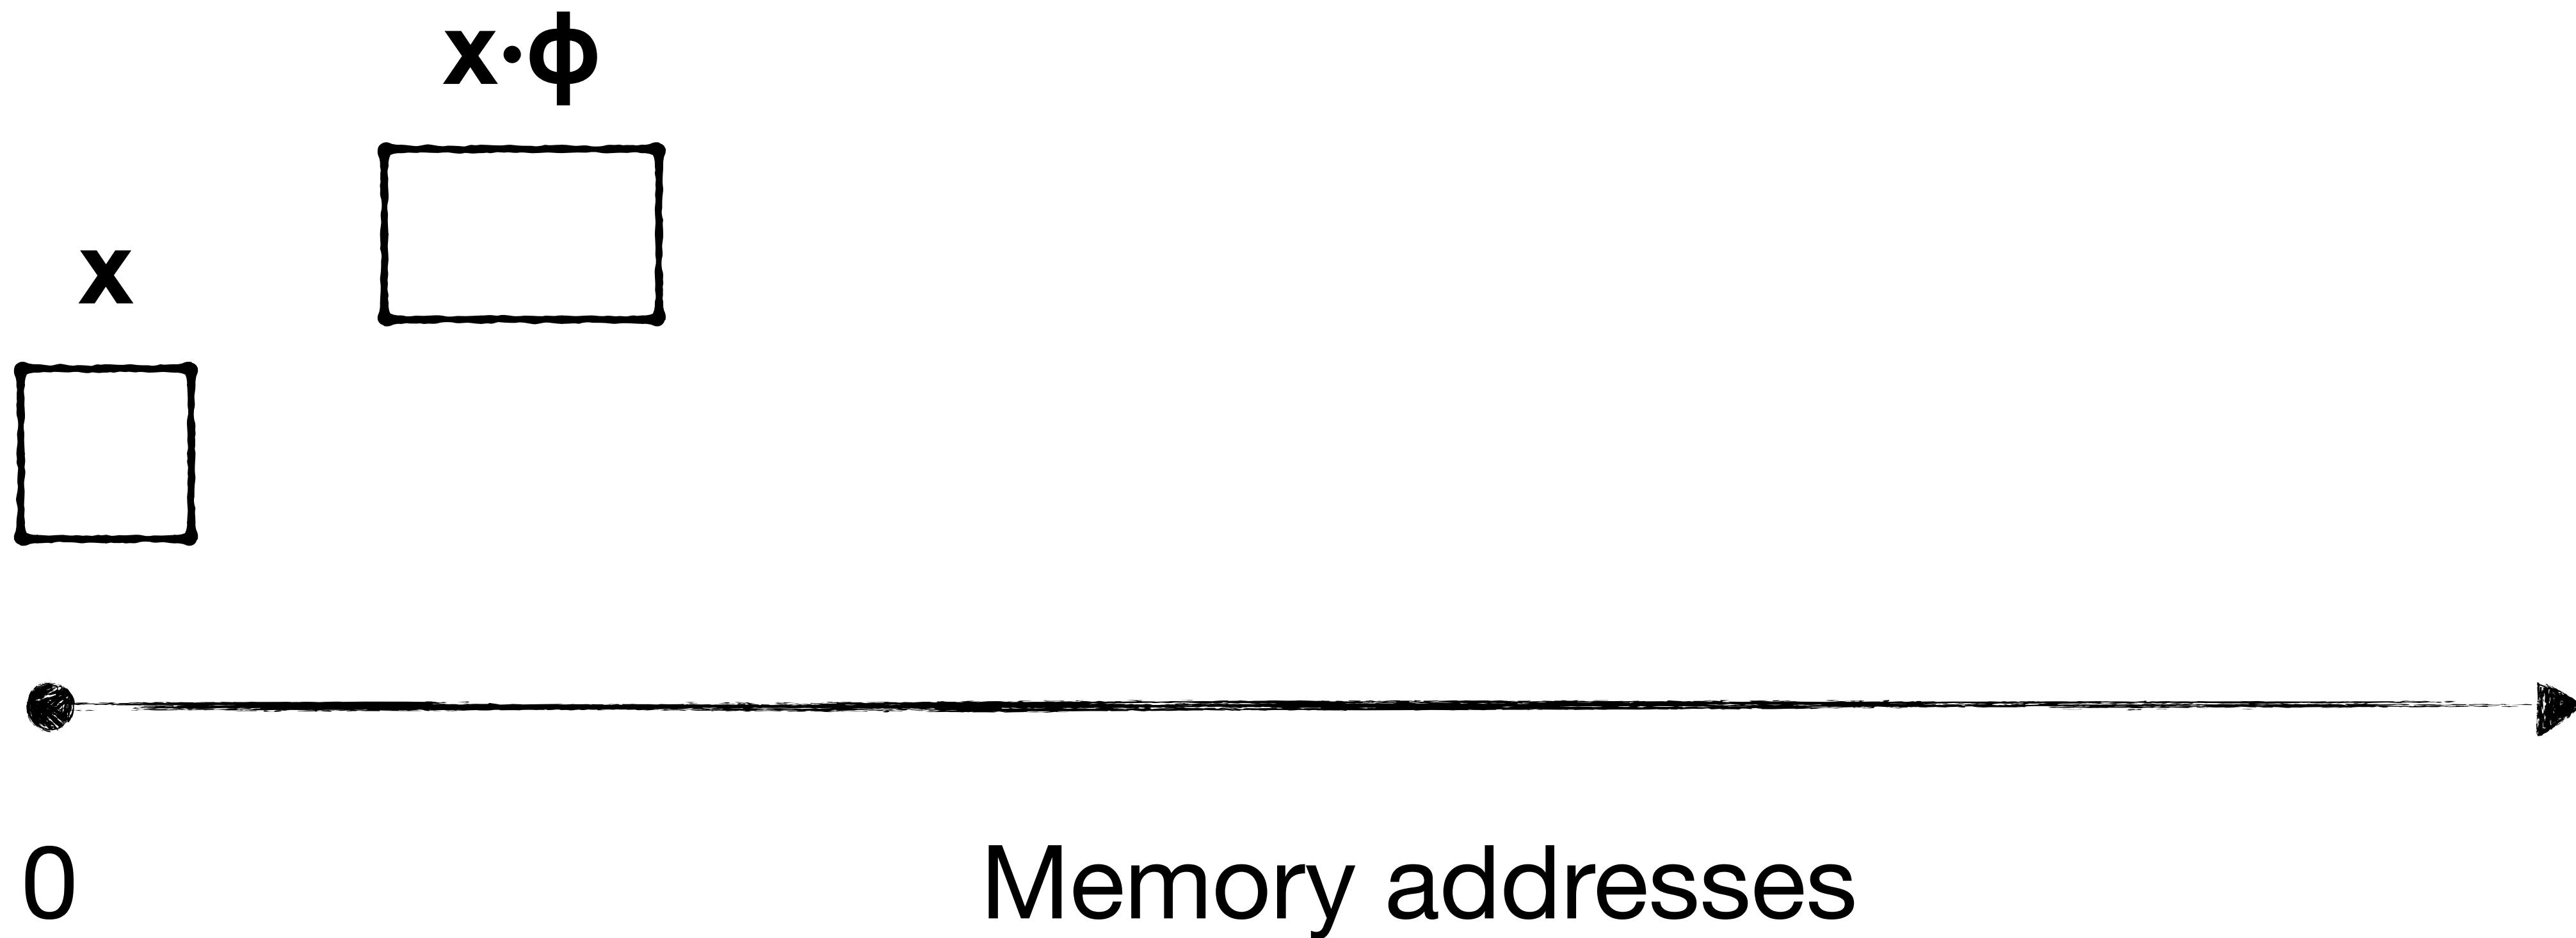
Satisfies both:


$$\text{Size}_{i-2} + \text{Size}_{i-1} = \text{Size}_i$$
$$\frac{\text{Size}_{i-1}}{\text{Size}_{i-2}} = \frac{\text{Size}_i}{\text{Size}_{i-1}} = G$$

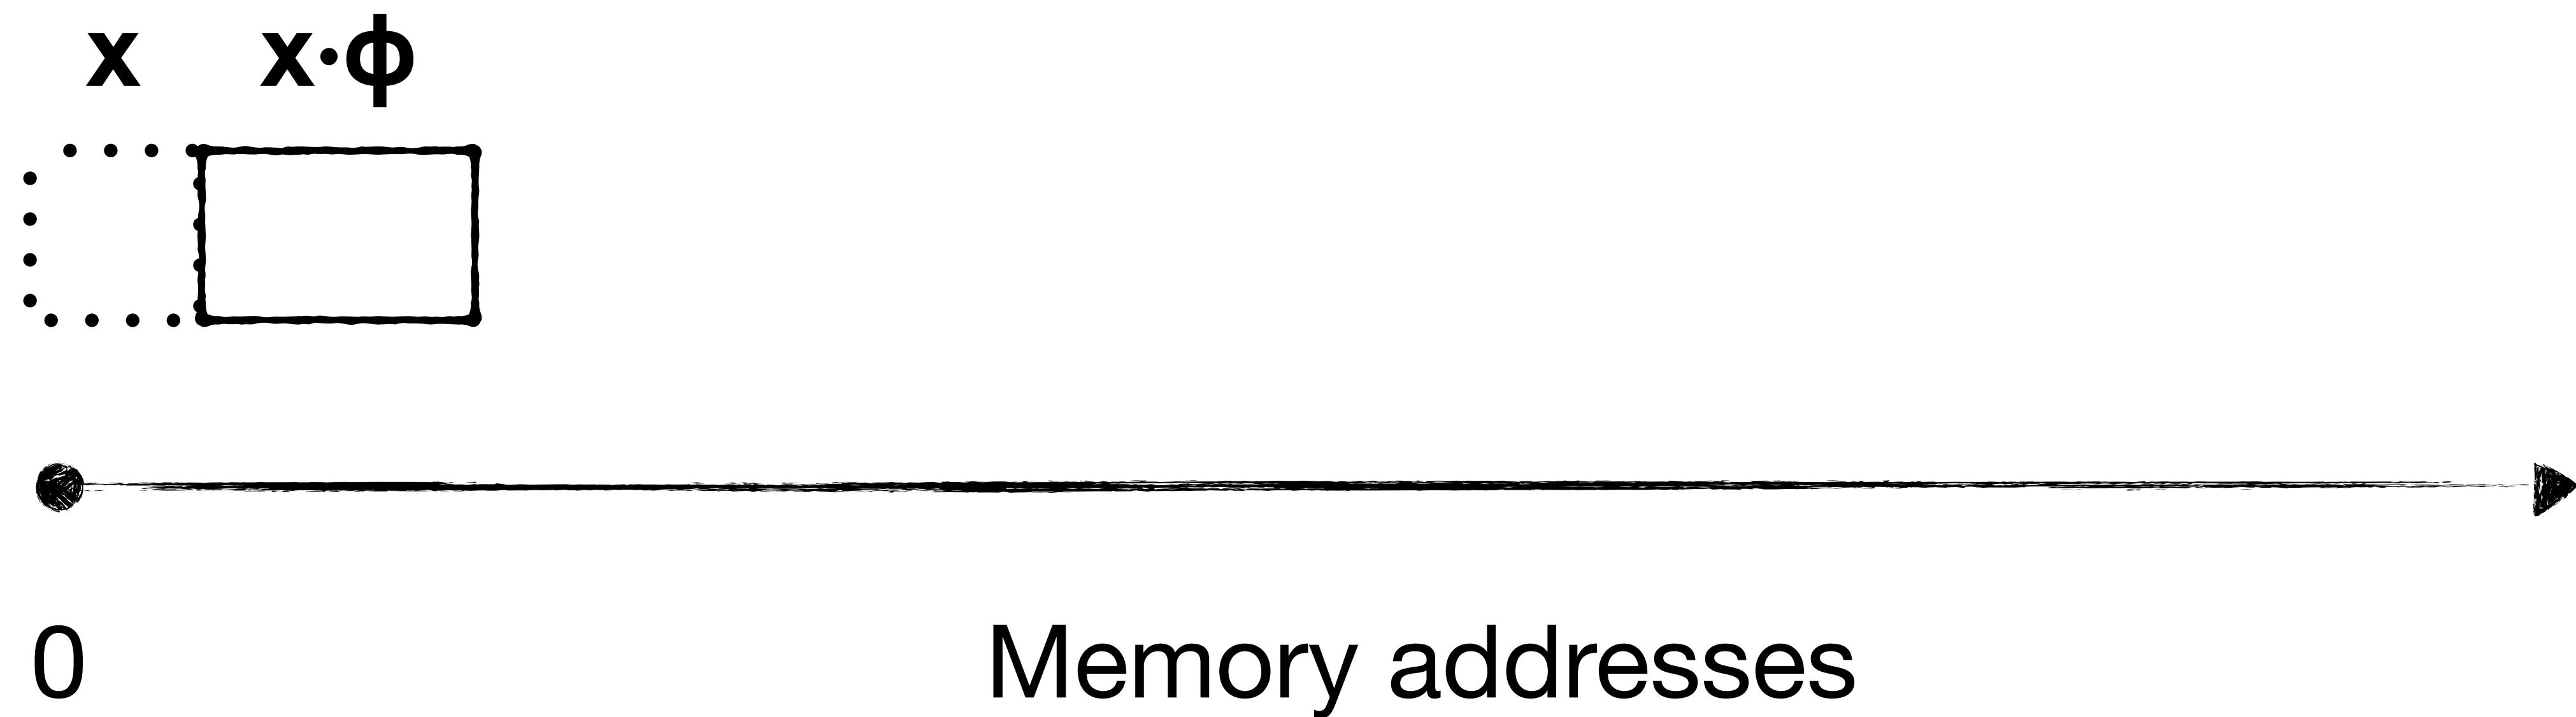
Expand Array by Golden Ratio ($G = \phi = 1.61\dots$)



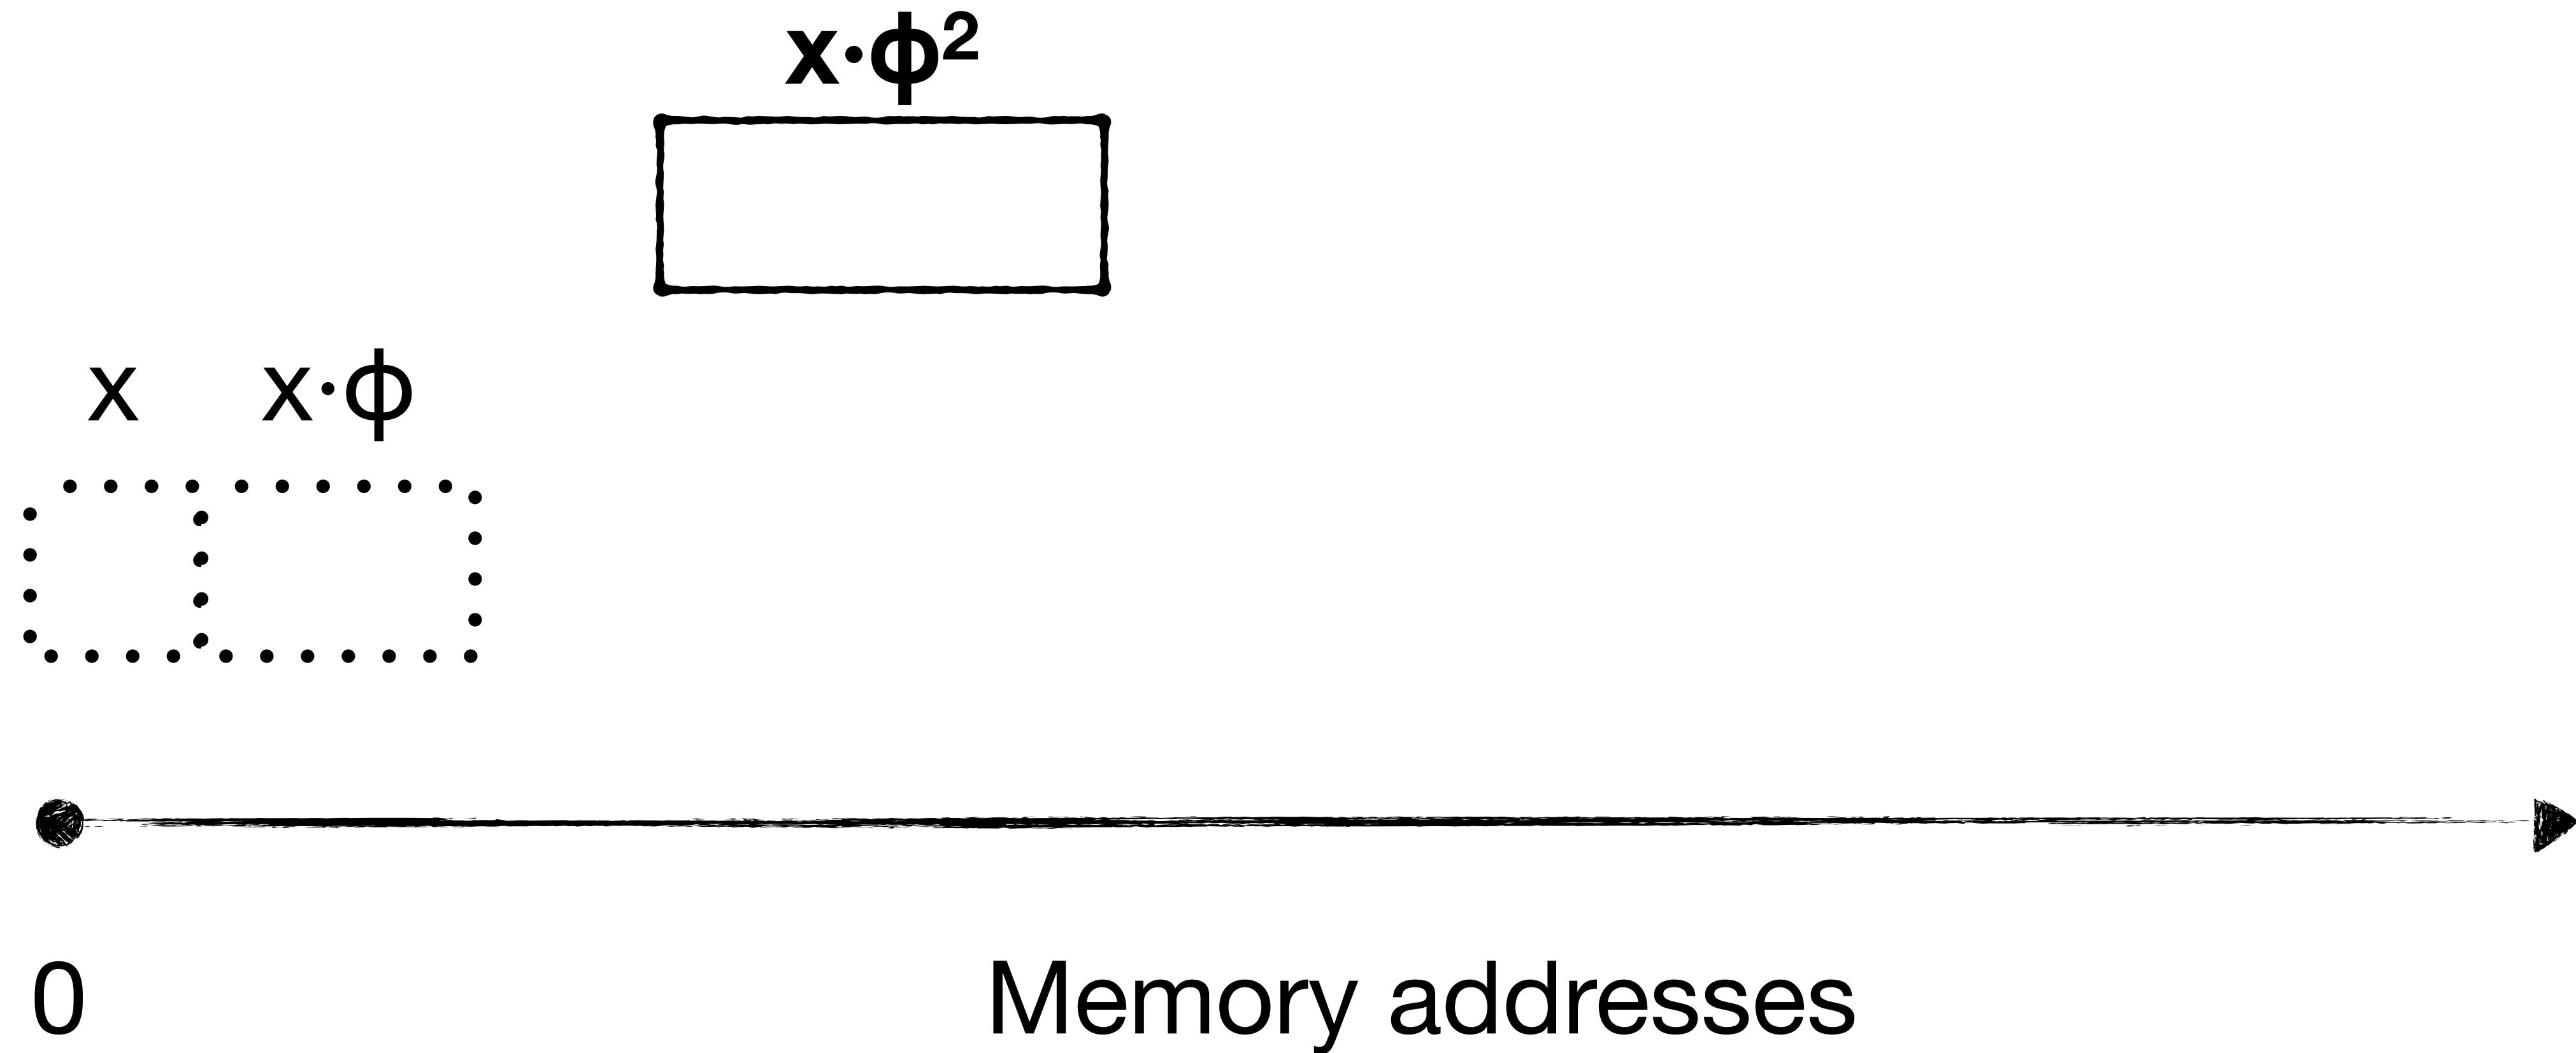
Expand Array by Golden Ratio ($G = \phi = 1.61\dots$)



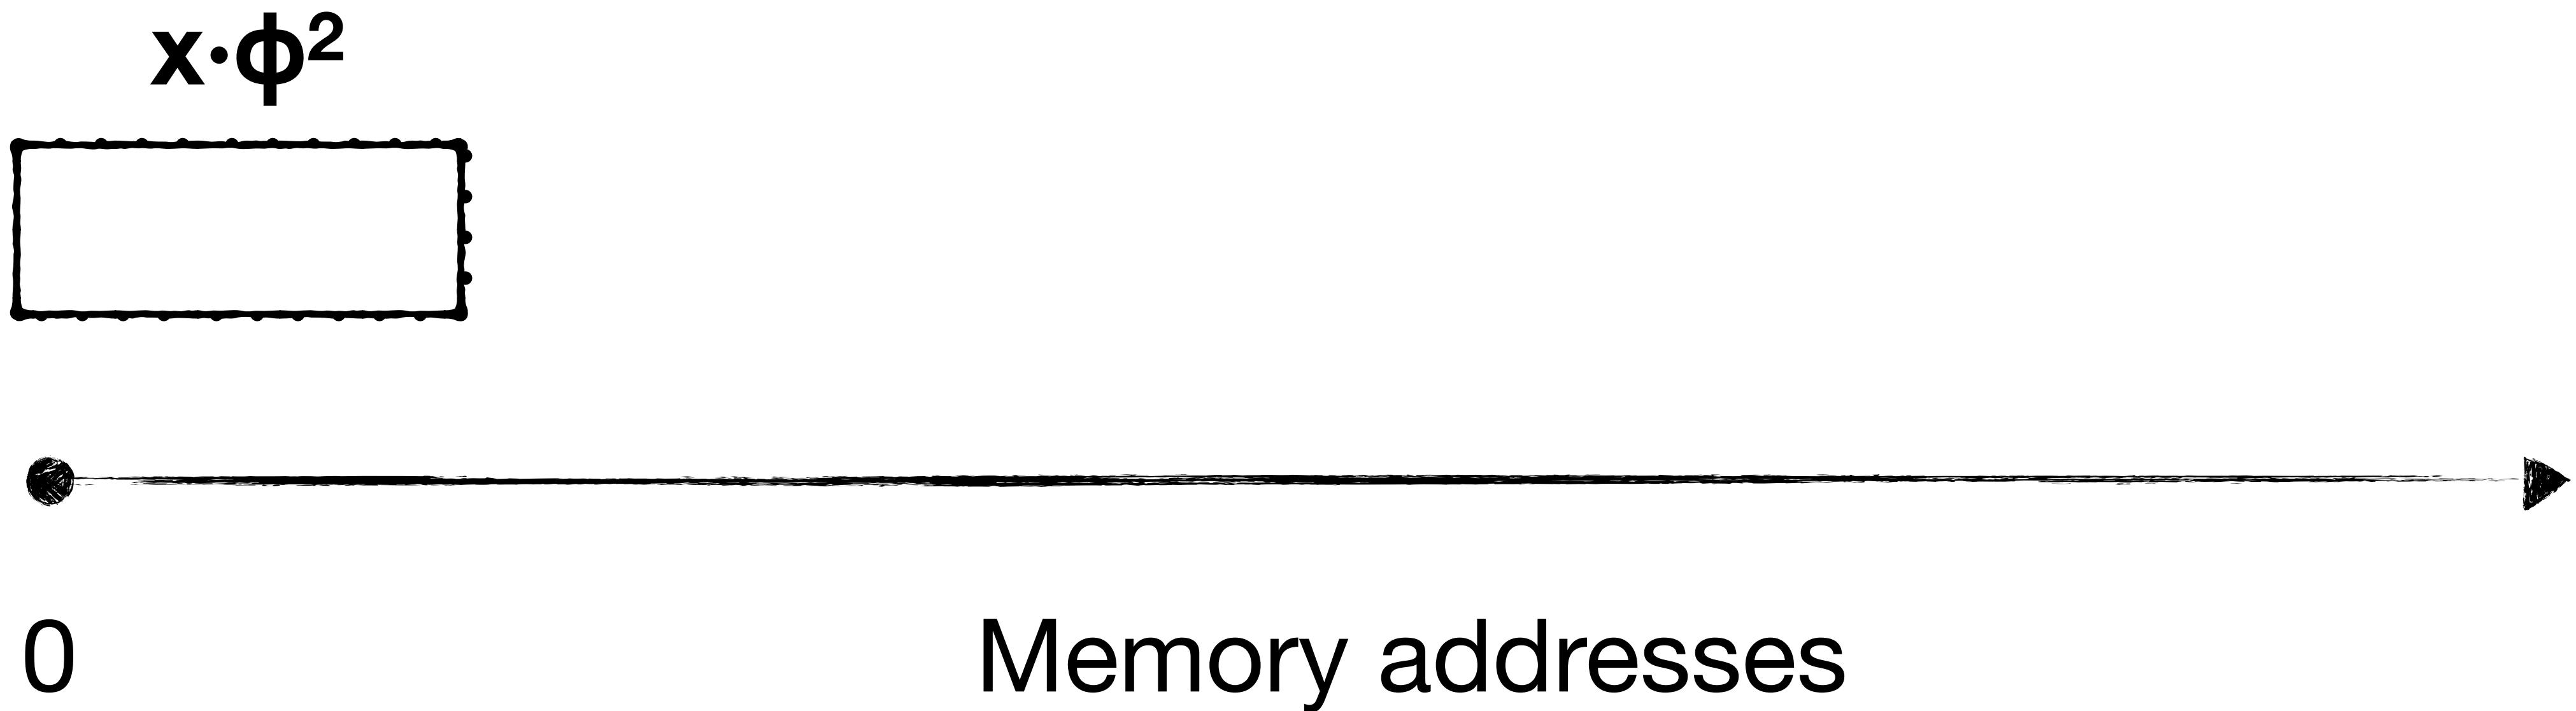
Expand Array by Golden Ratio ($G = \phi = 1.61\dots$)



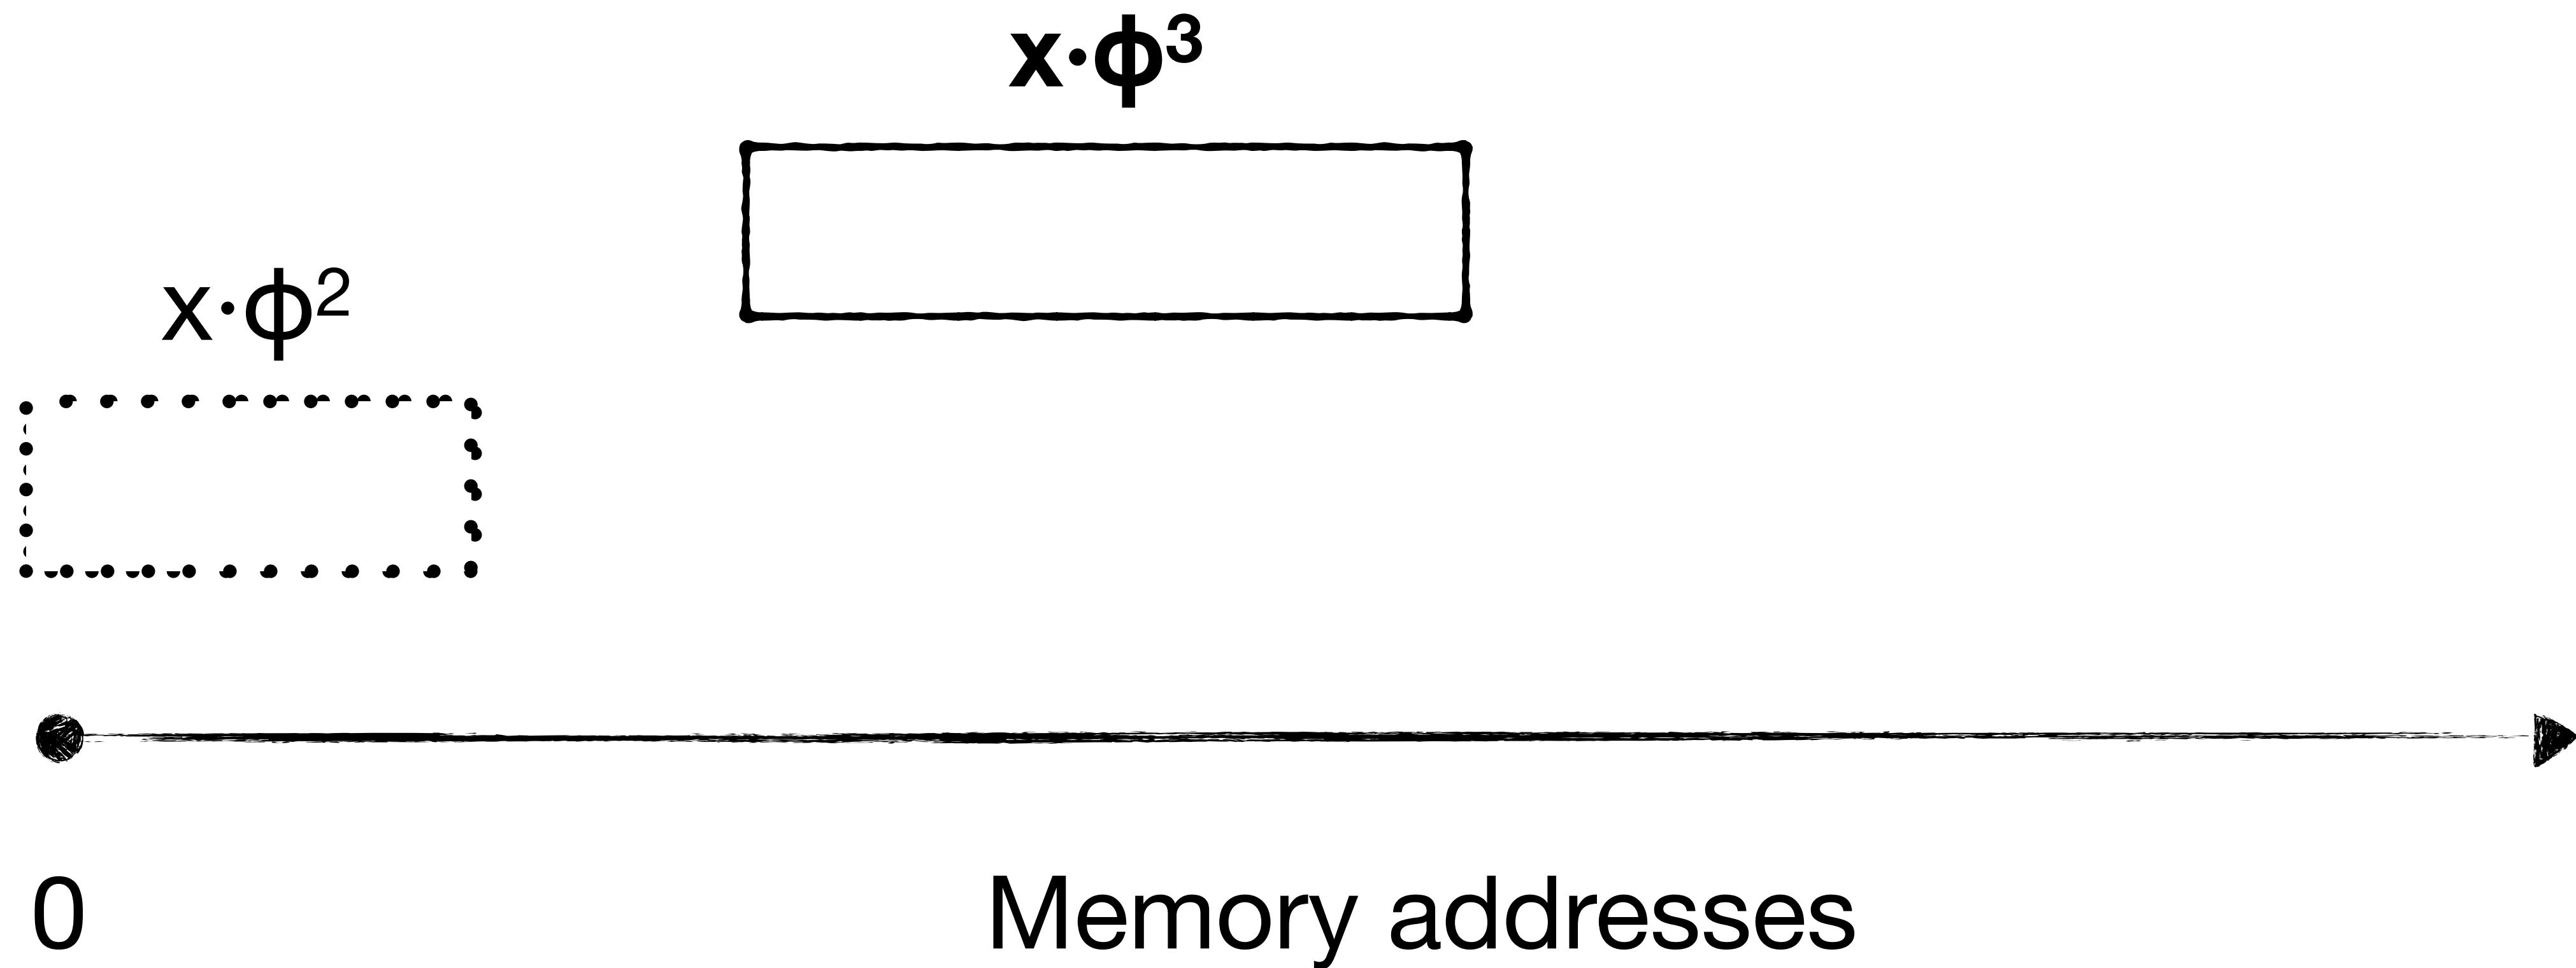
Expand Array by Golden Ratio ($G = \phi = 1.61\dots$)



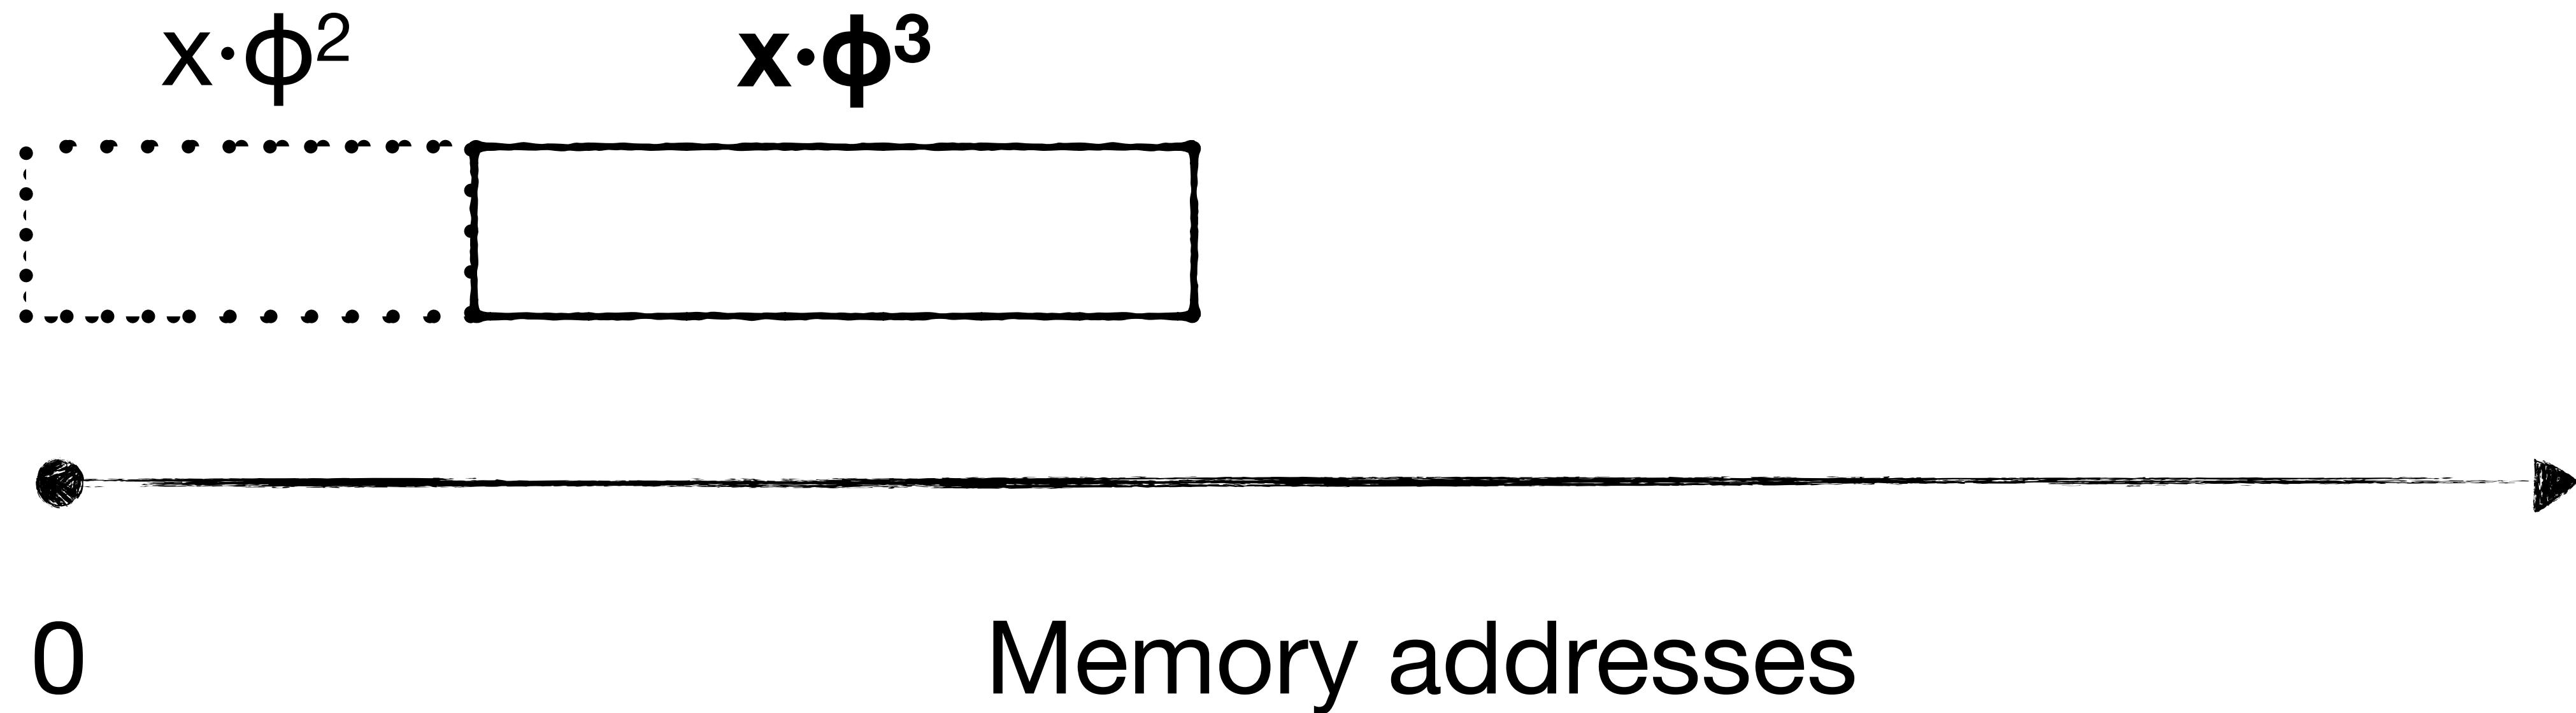
Expand Array by Golden Ratio ($G = \phi = 1.61\dots$)



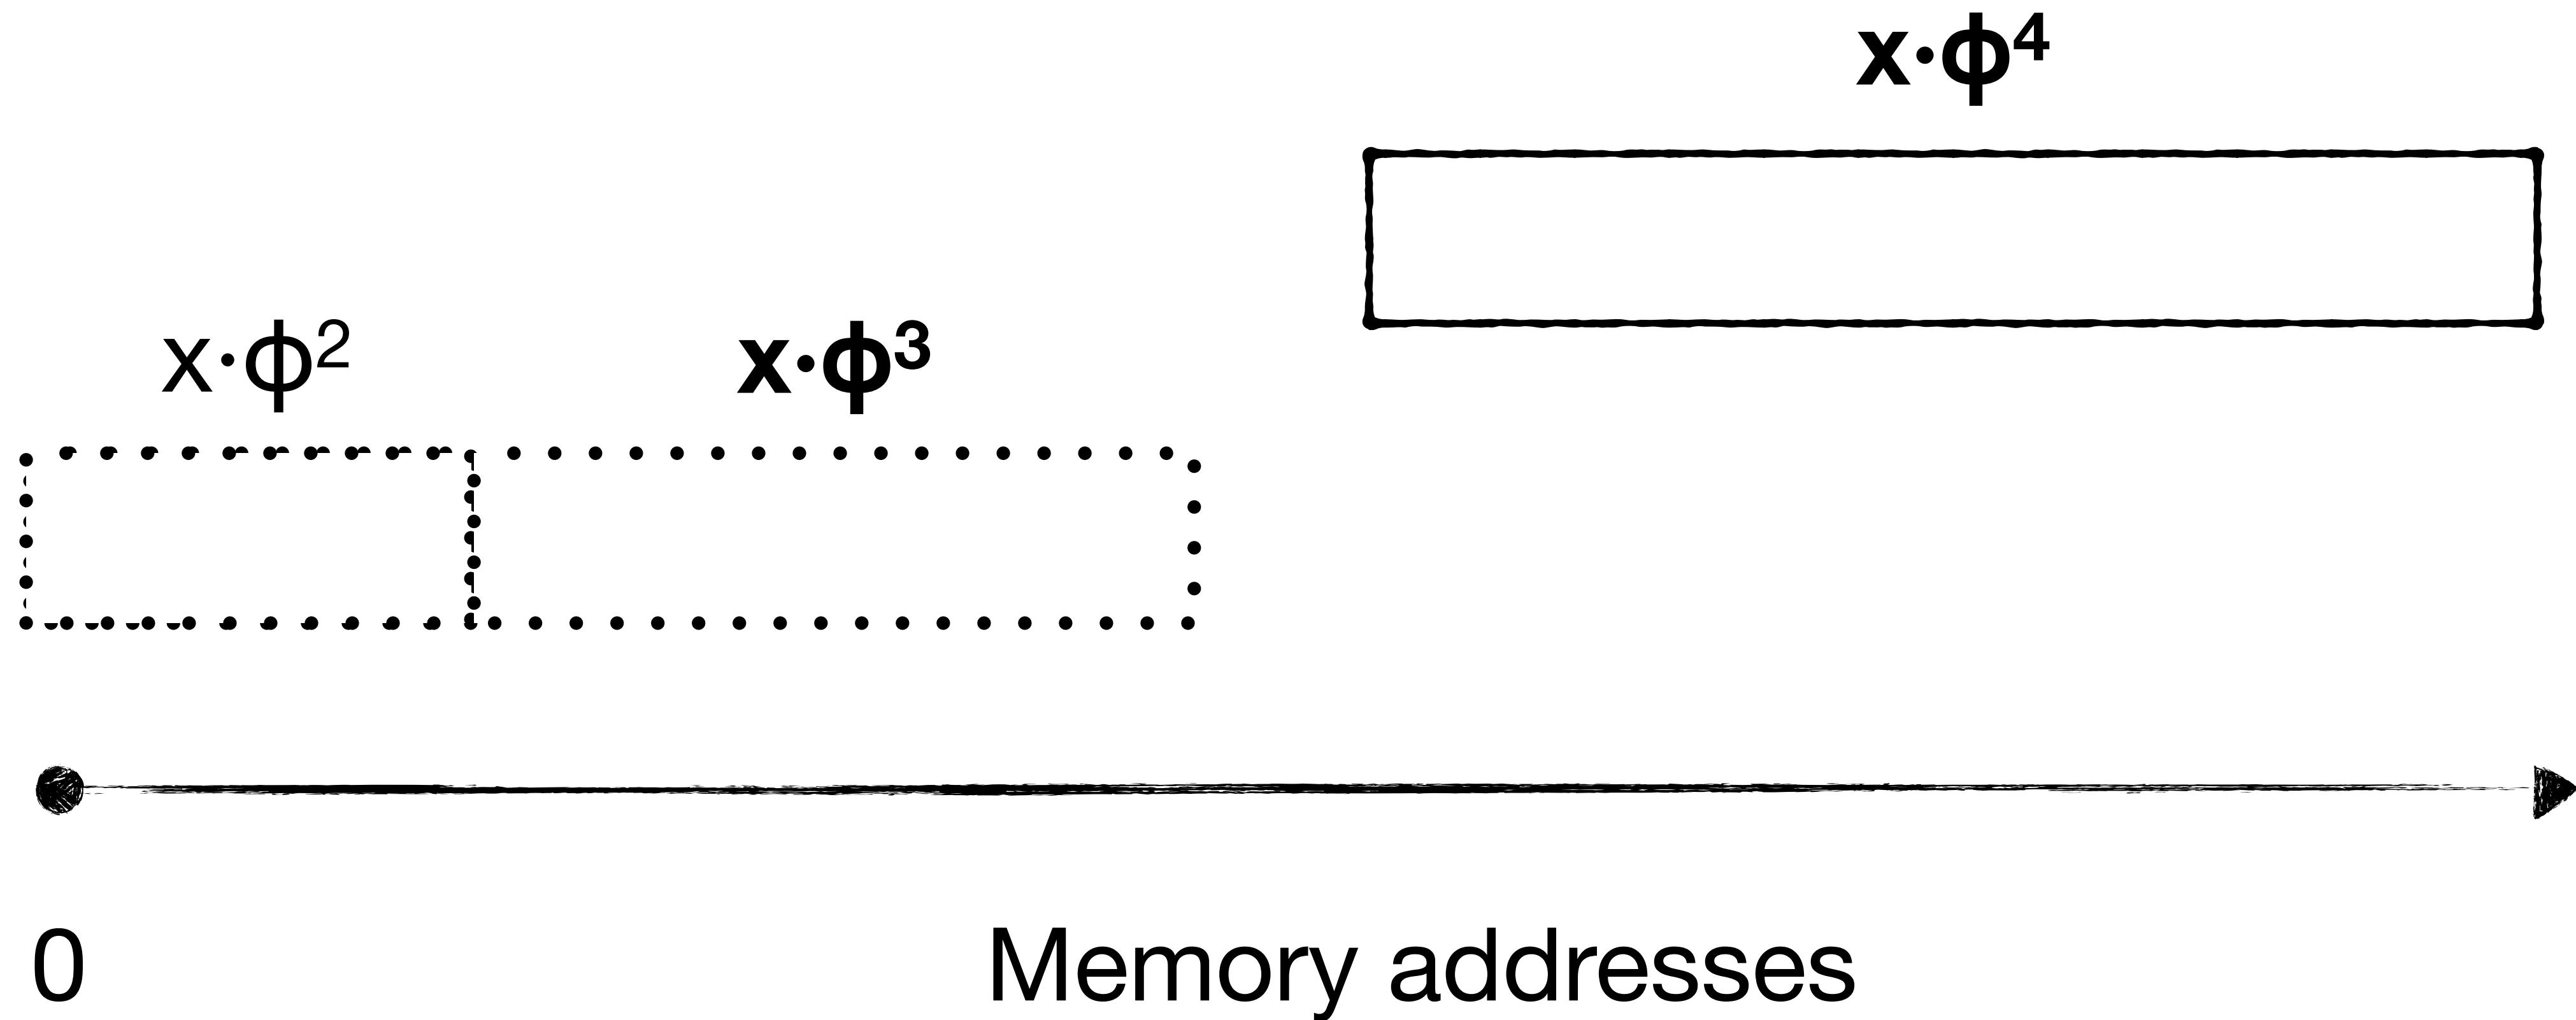
Expand Array by Golden Ratio ($G = \phi = 1.61\dots$)



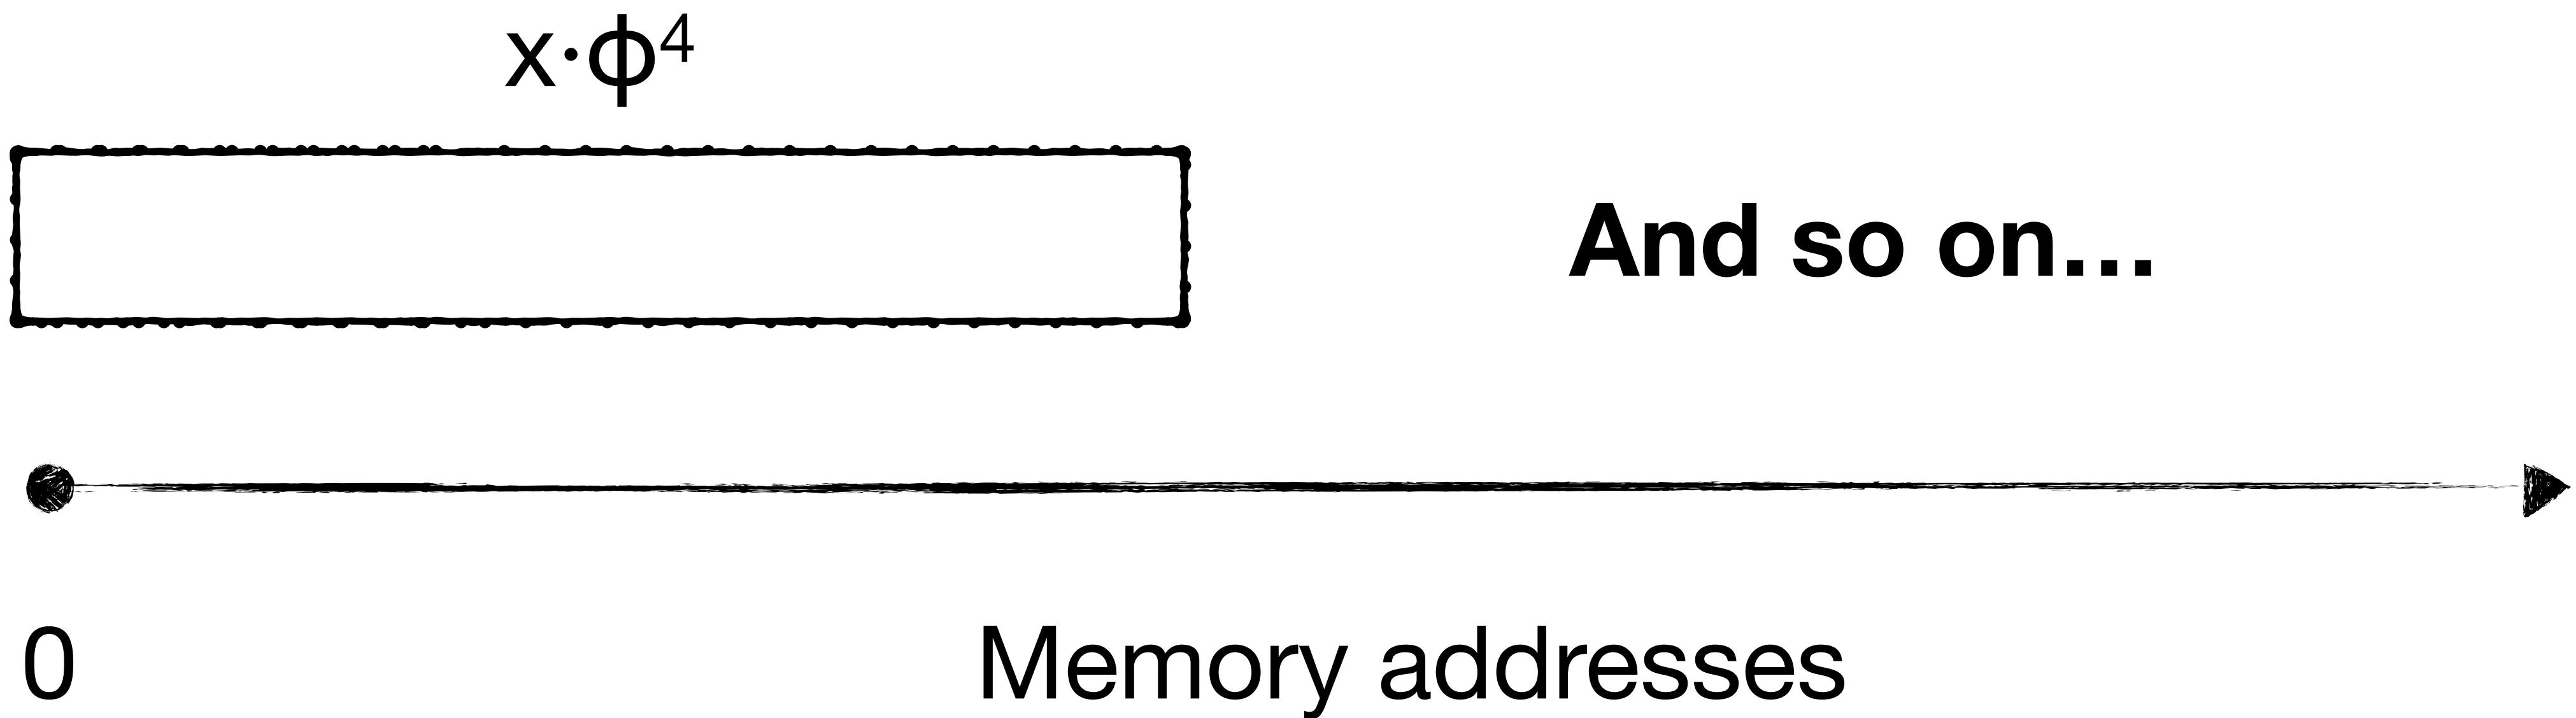
Expand Array by Golden Ratio ($G = \phi = 1.61\dots$)



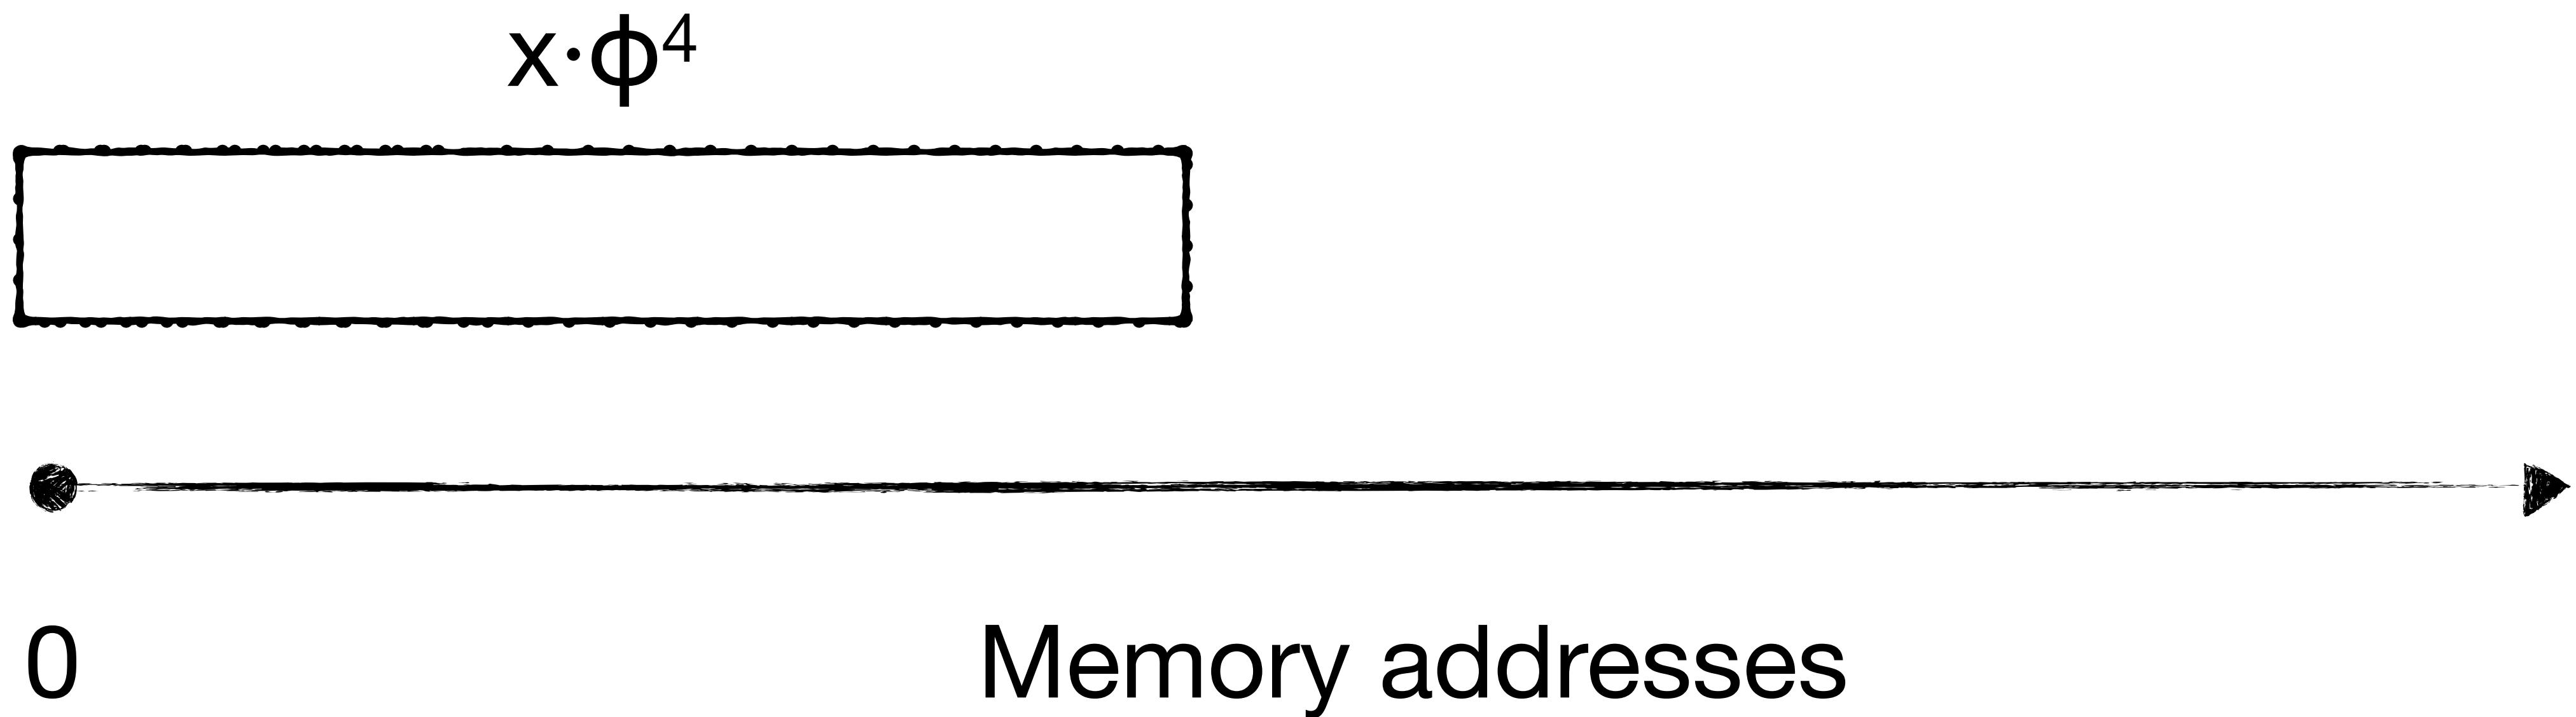
Expand Array by Golden Ratio ($G = \phi = 1.61\dots$)



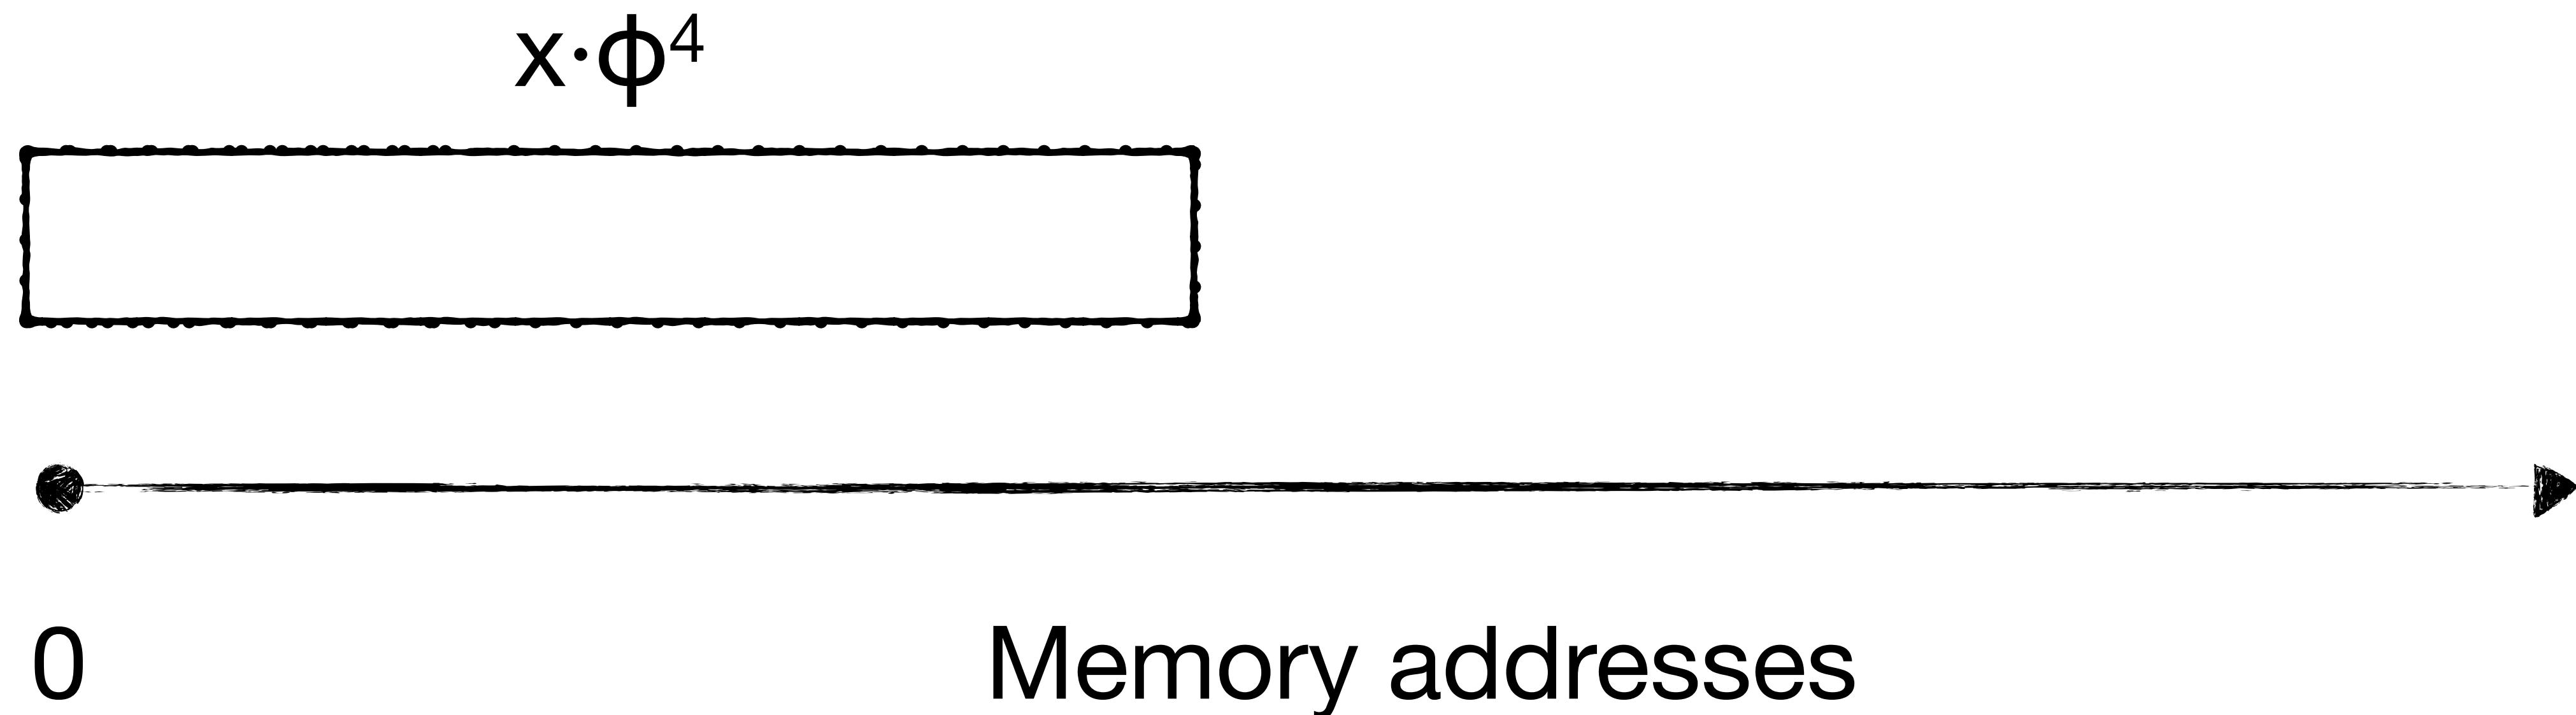
Expand Array by Golden Ratio ($G = \phi = 1.61\dots$)



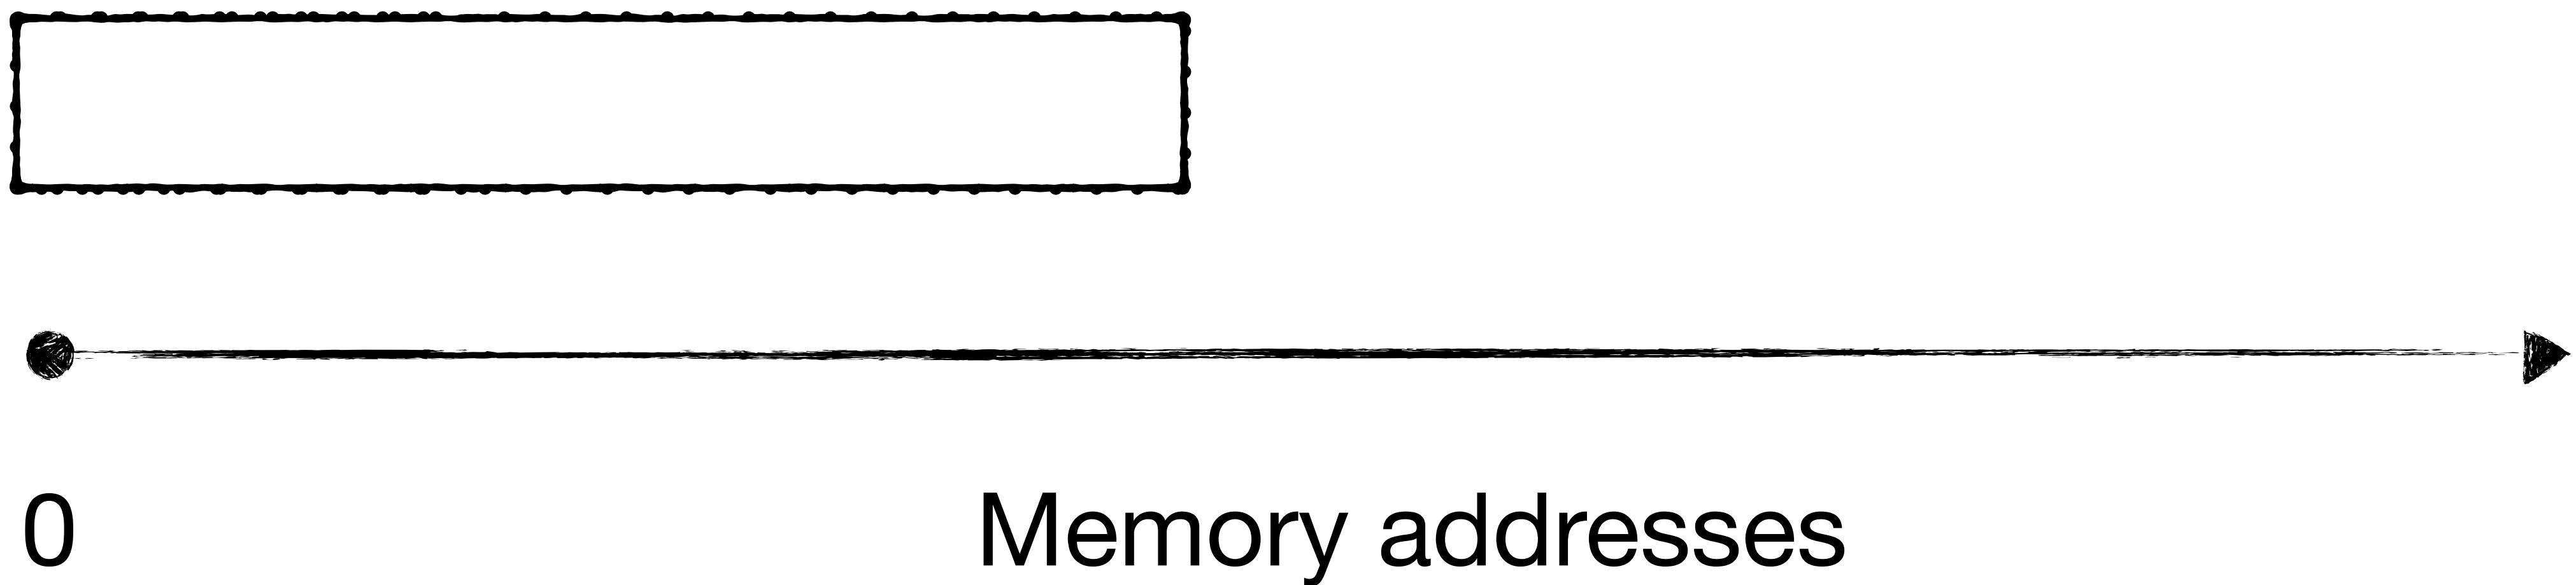
Write-Amplification



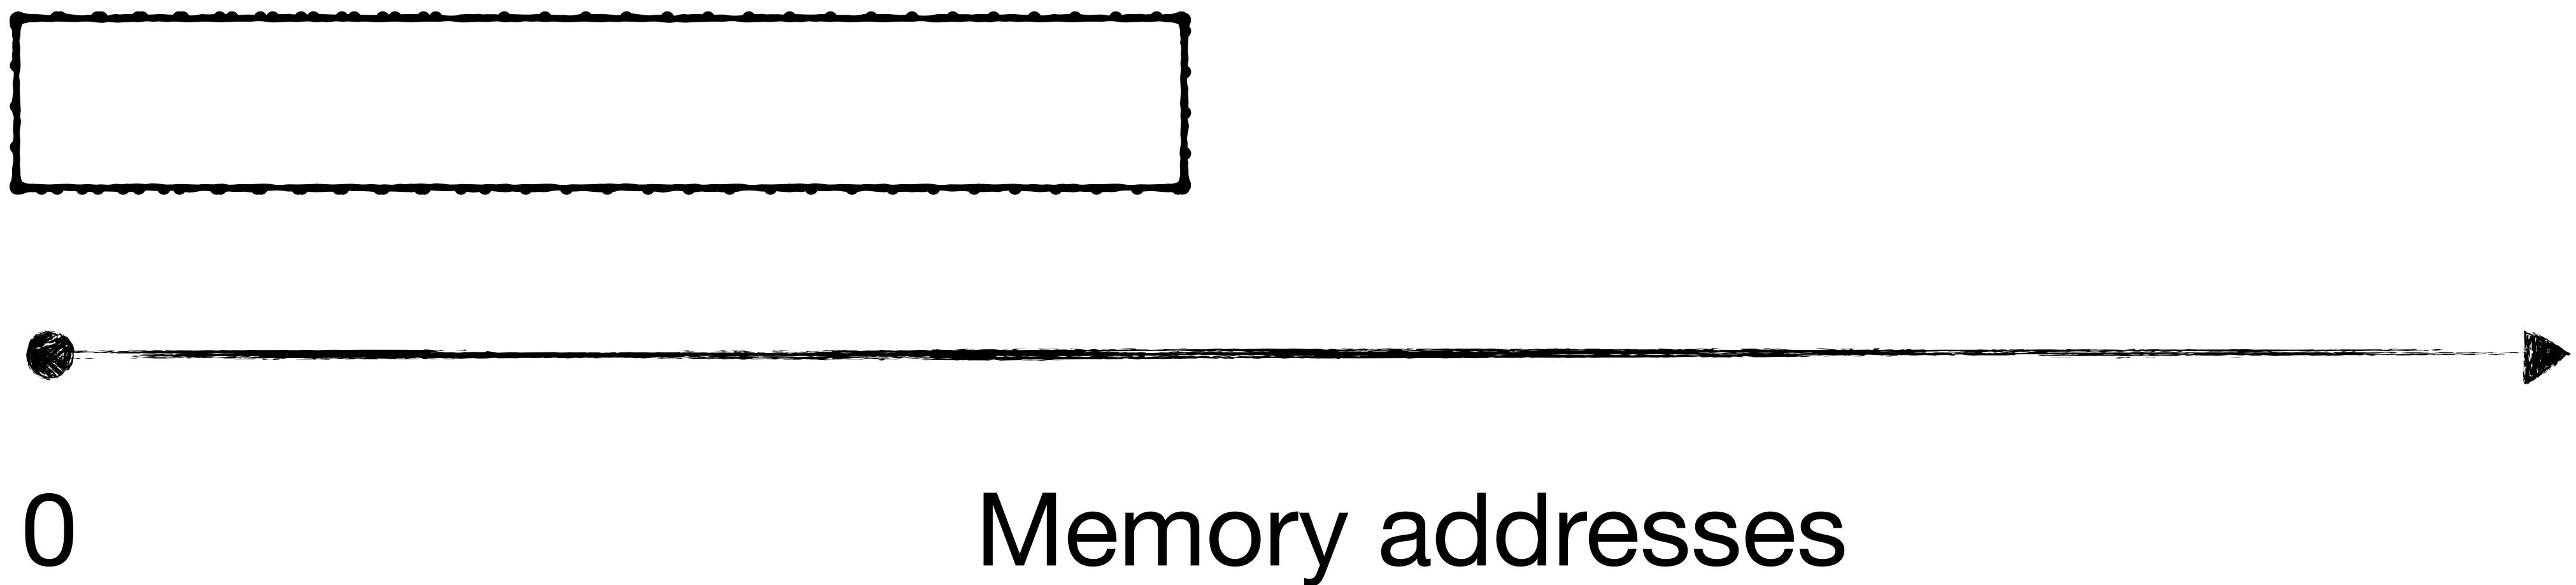
$$\text{Write-Amplification} = \frac{G}{G - 1}$$



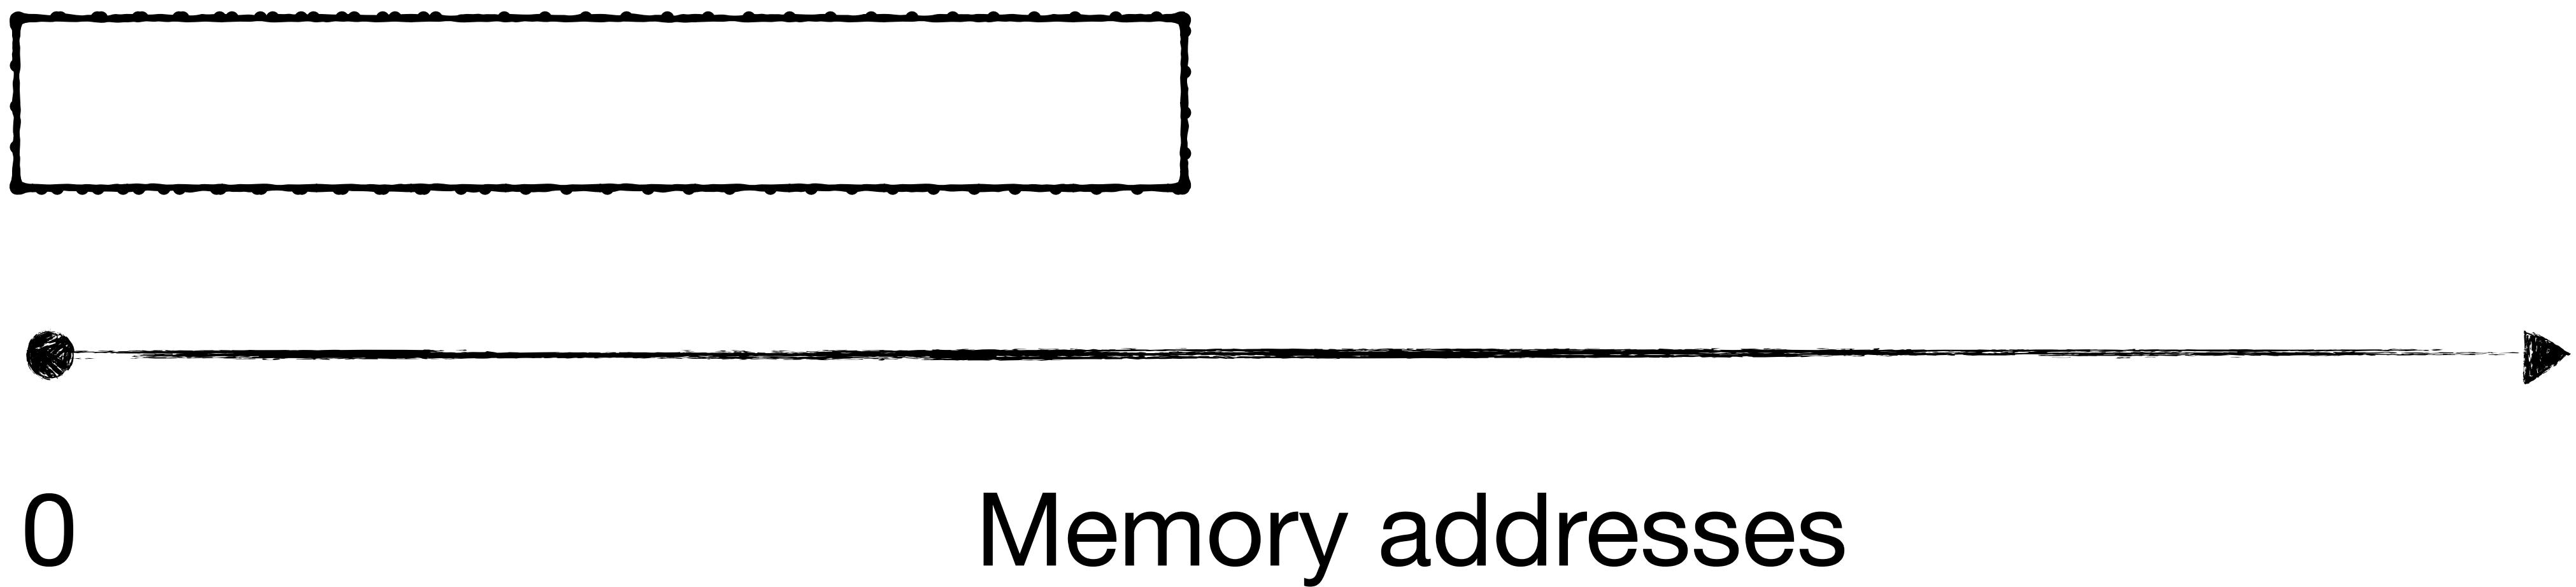
$$\text{Write-Amplification} = \frac{G}{G - 1} = \frac{\phi}{\phi - 1}$$



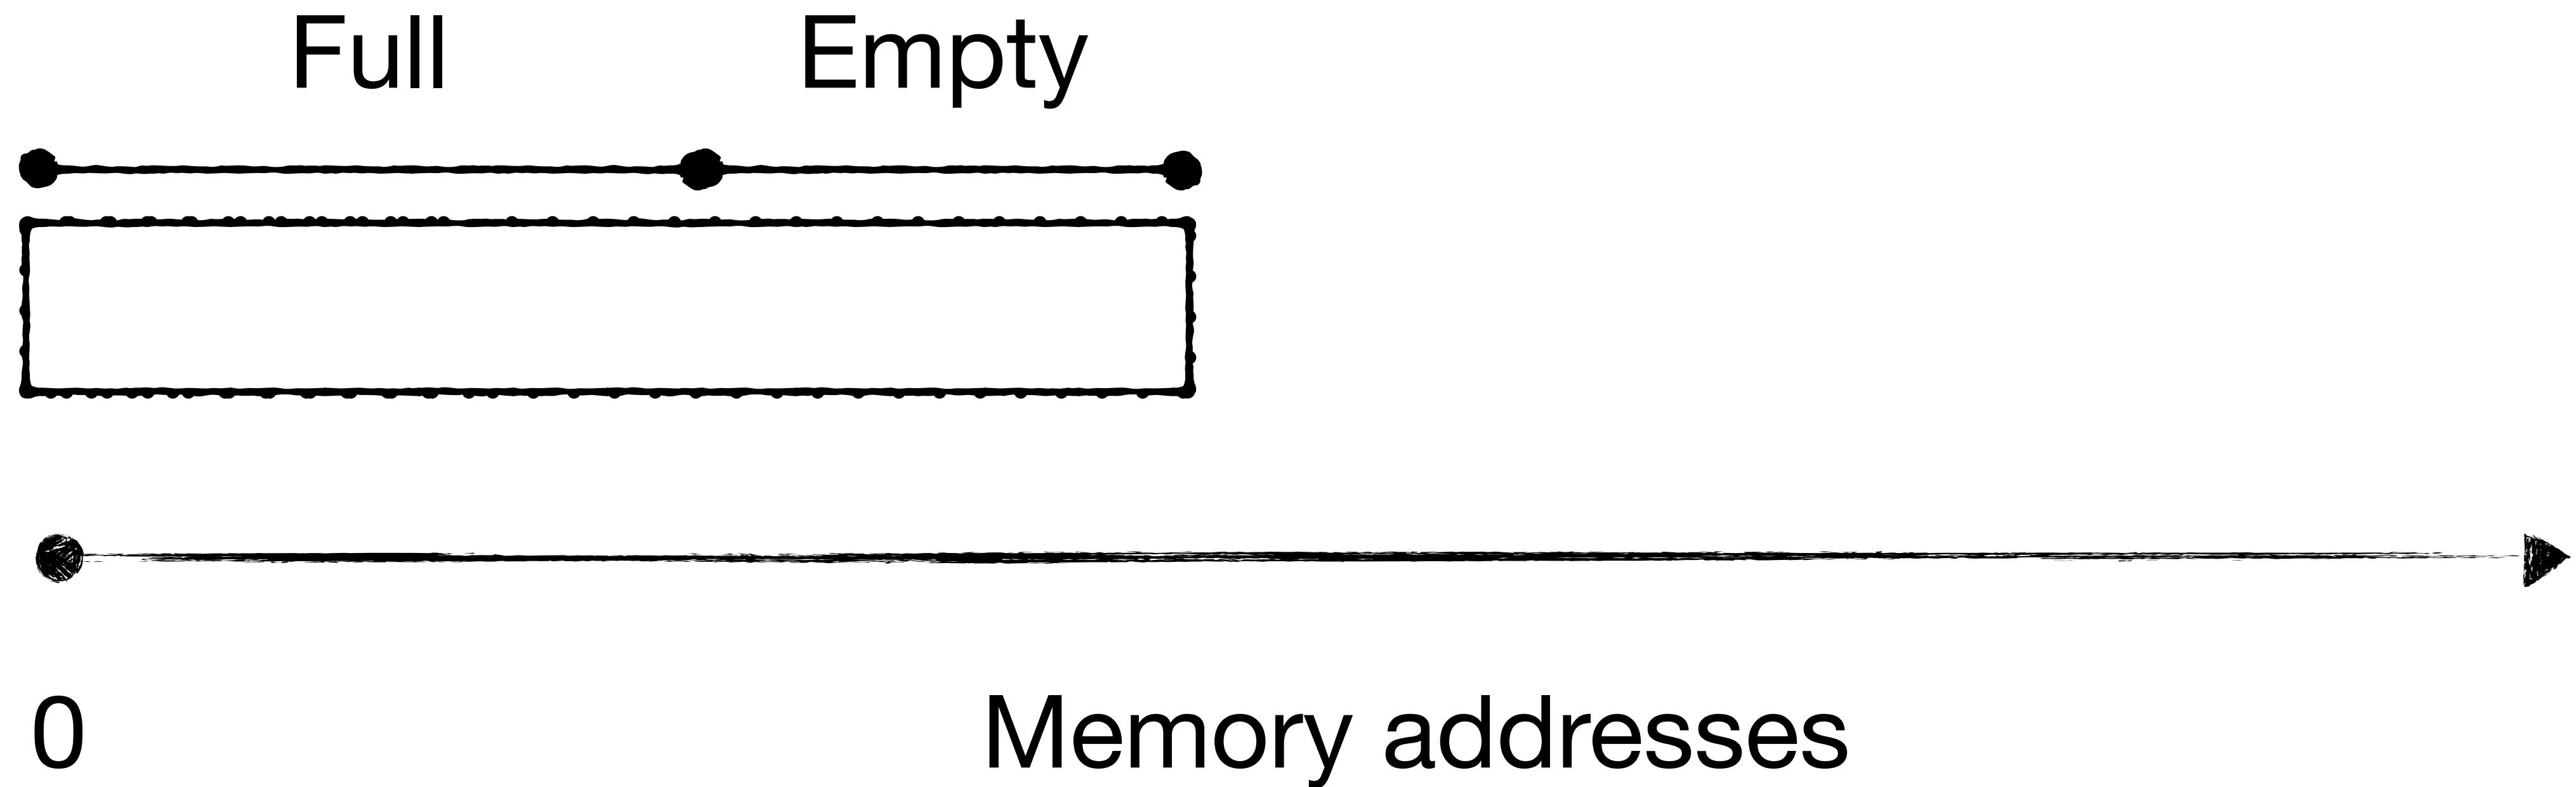
$$\text{Write-Amplification} = \frac{G}{G - 1} = \frac{\phi}{\phi - 1} = \phi + 1$$



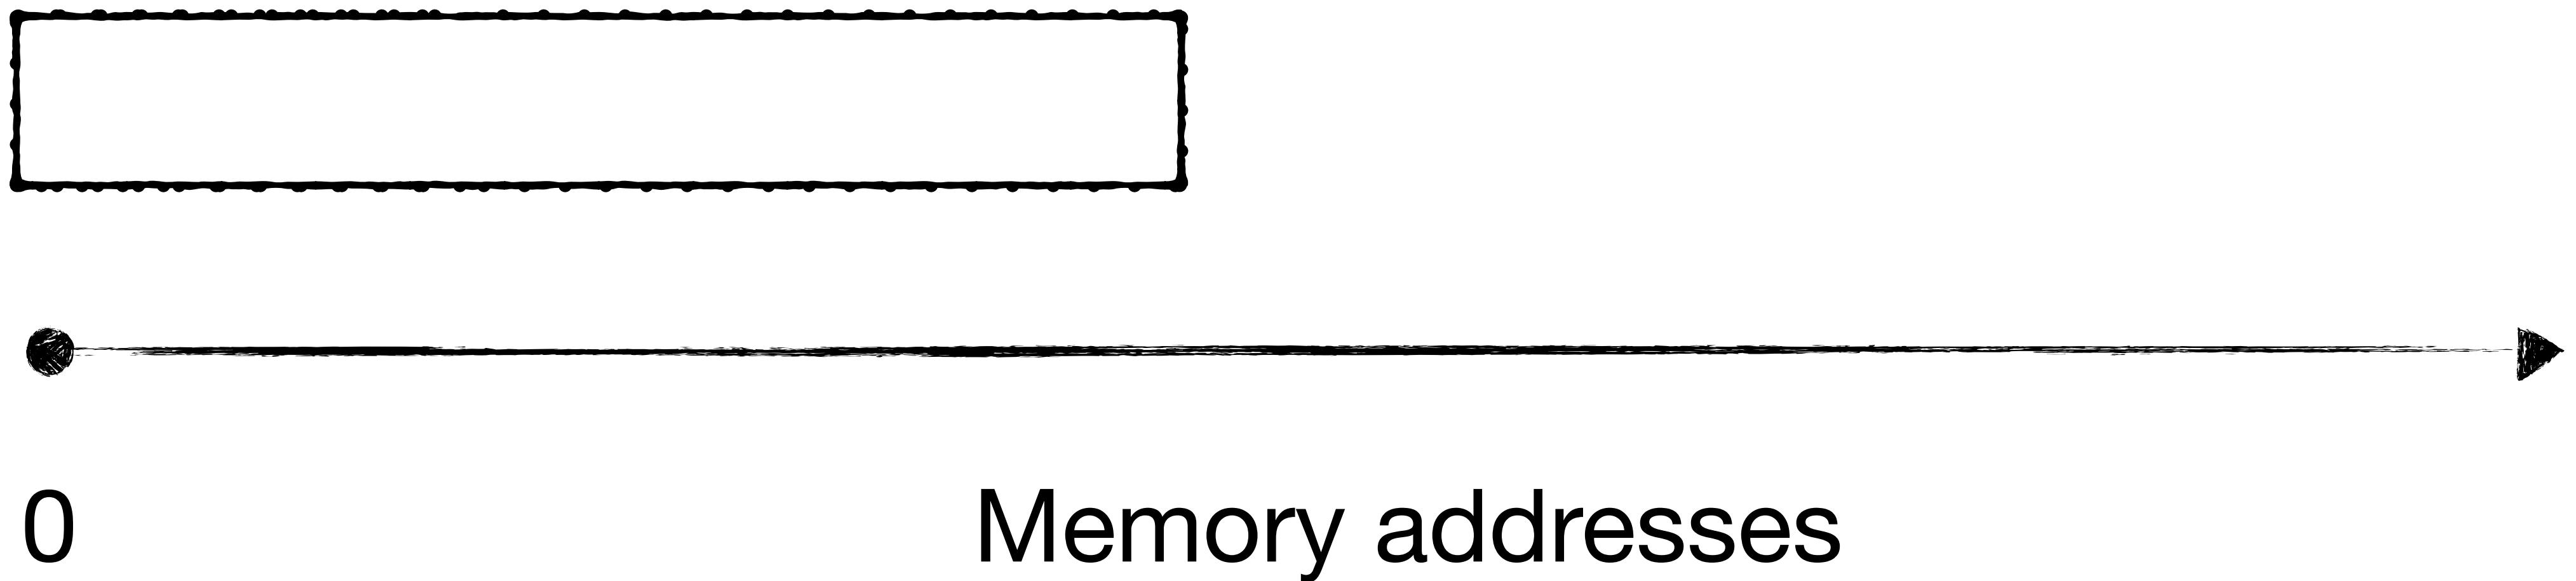
Space-Amplification?



Space-Amplification?



$$\text{Space-Amplification} = \frac{\text{Full} + \text{Empty}}{\text{Full}} = G = \phi$$

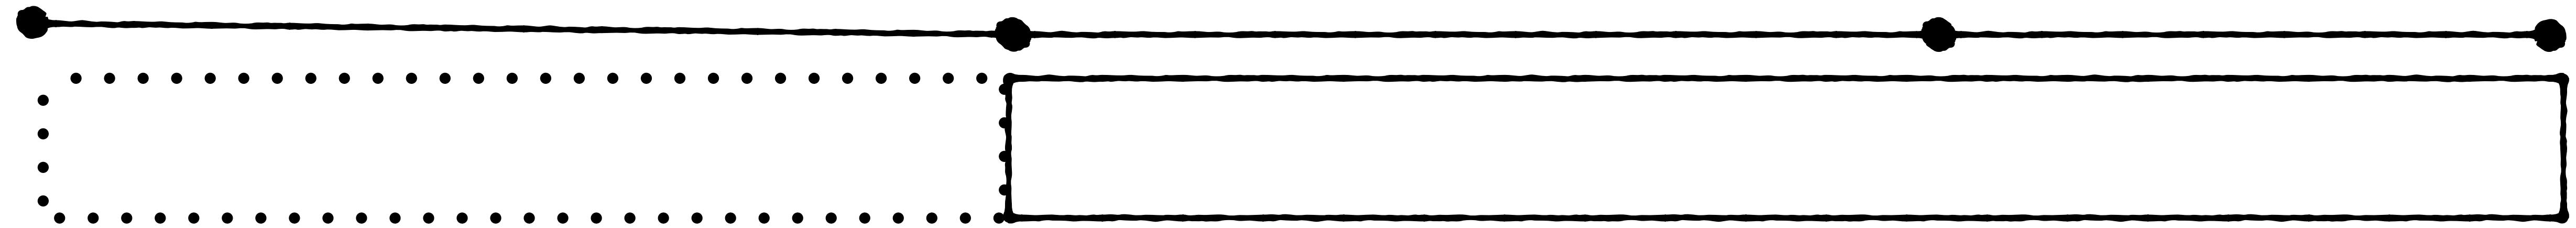


Max Space-Amp?

Deallocated

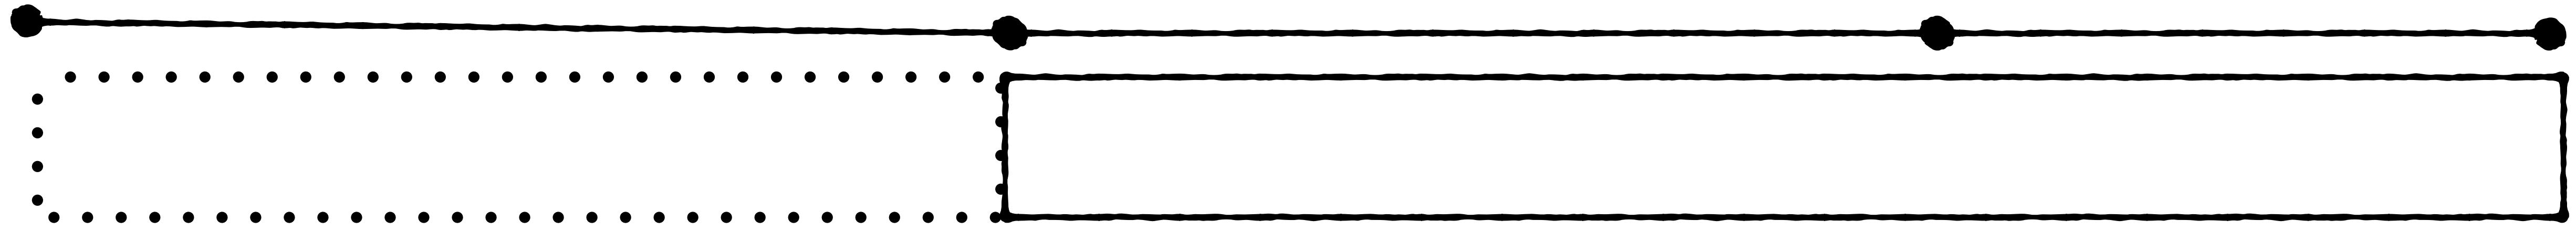
Full

Empty

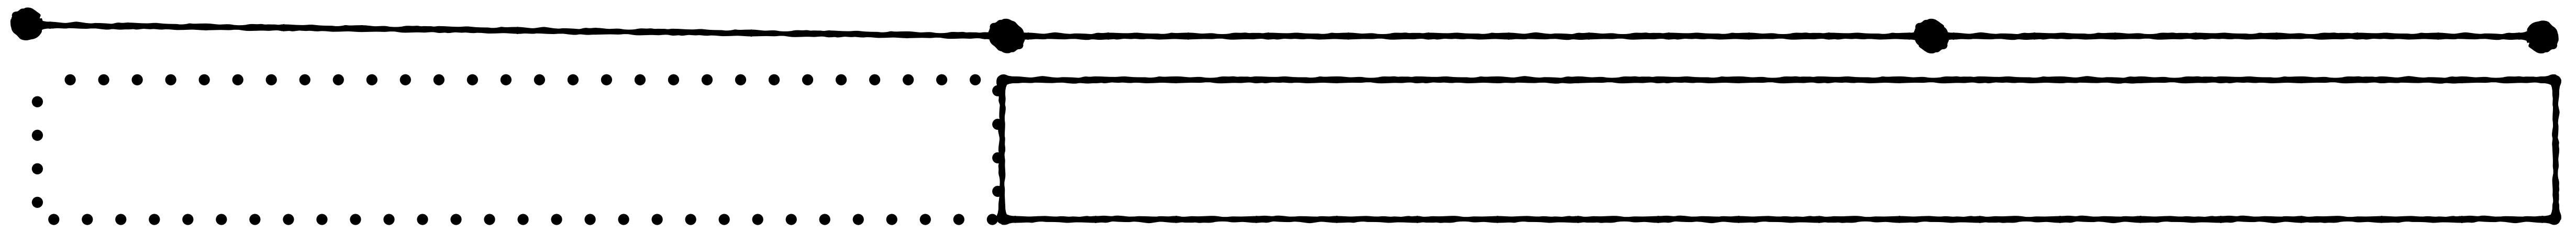


For $G < \phi$

$$\text{Max Space-Amp} = \frac{\text{Deallocated} + \text{Full} + \text{Empty}}{\text{Full}} = 1 + G$$



$$\text{Max Space-Amp} = 1 + \phi = 2.61$$



Write-amp

$G > \phi$

$$\frac{G}{G - 1}$$

$G < \phi$

$$\frac{G}{G - 1}$$

Space-amp

$$\frac{G}{G - 1} + G$$

Alternates
G to G+1

In the wild

Implementation	Growth factor
Java ArrayList	1.5
Python PyListObject	~1.125
Microsoft Visual C++ 2013	1.5
G++ 5.2.0	2
Clang 3.6	2
Facebook folly/FBVector	1.5
Rust Vec	2
Go slices	between 1.25 and 2
Nim sequences	2
SBCL (Common Lisp) vectors	2
C# (.NET 8) List	2

Source: https://en.wikipedia.org/wiki/Dynamic_array#Growth_factor

Facebook folly/FBVector

<https://github.com/facebook/folly/blob/main/folly/docs/FBVector.md>

Real-world discussion of these issues

Facebook folly/FBVector

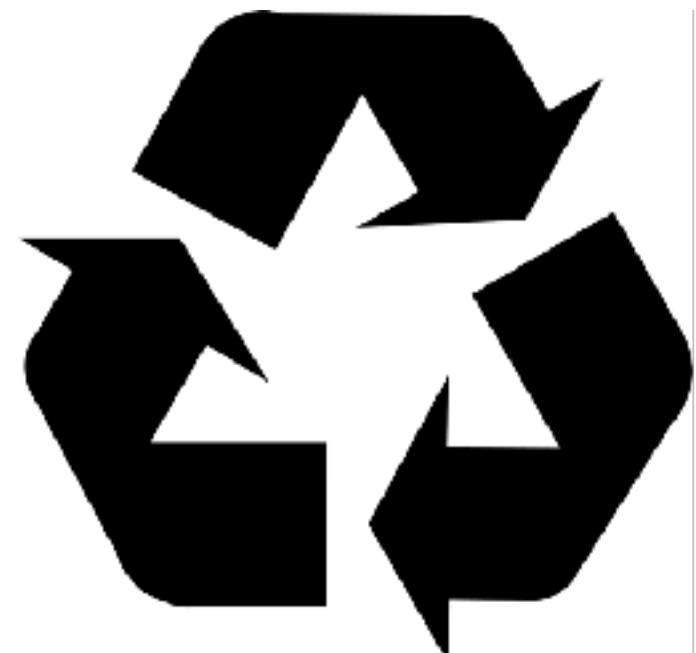
<https://github.com/facebook/folly/blob/main/folly/docs/FBVector.md>

Real-world discussion of these issues

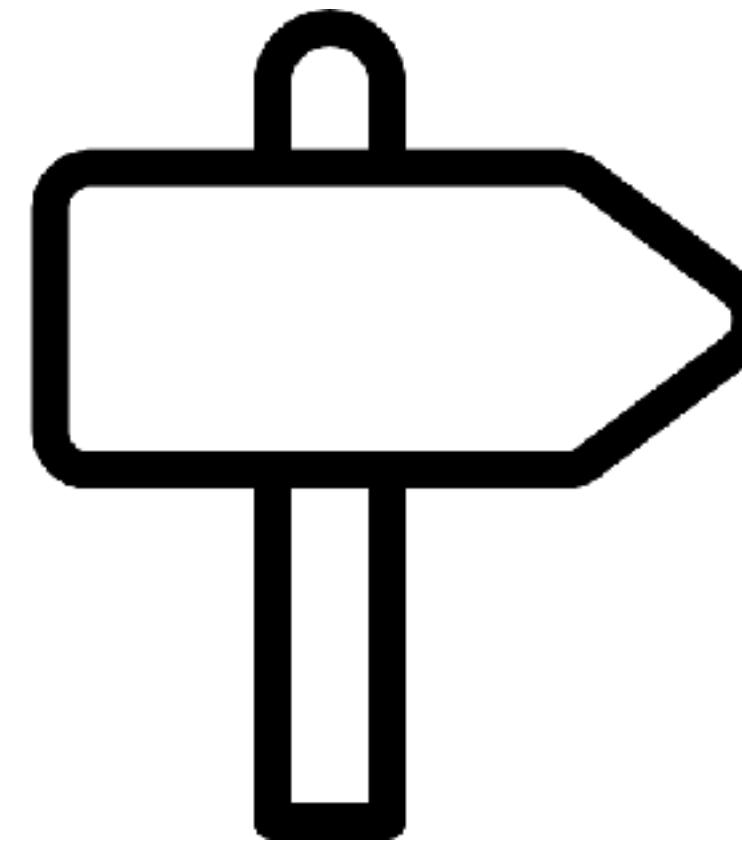
Note that Facebook also makes their own memory allocator, so with full control of the stack this can be more effective.

And now to new stuff

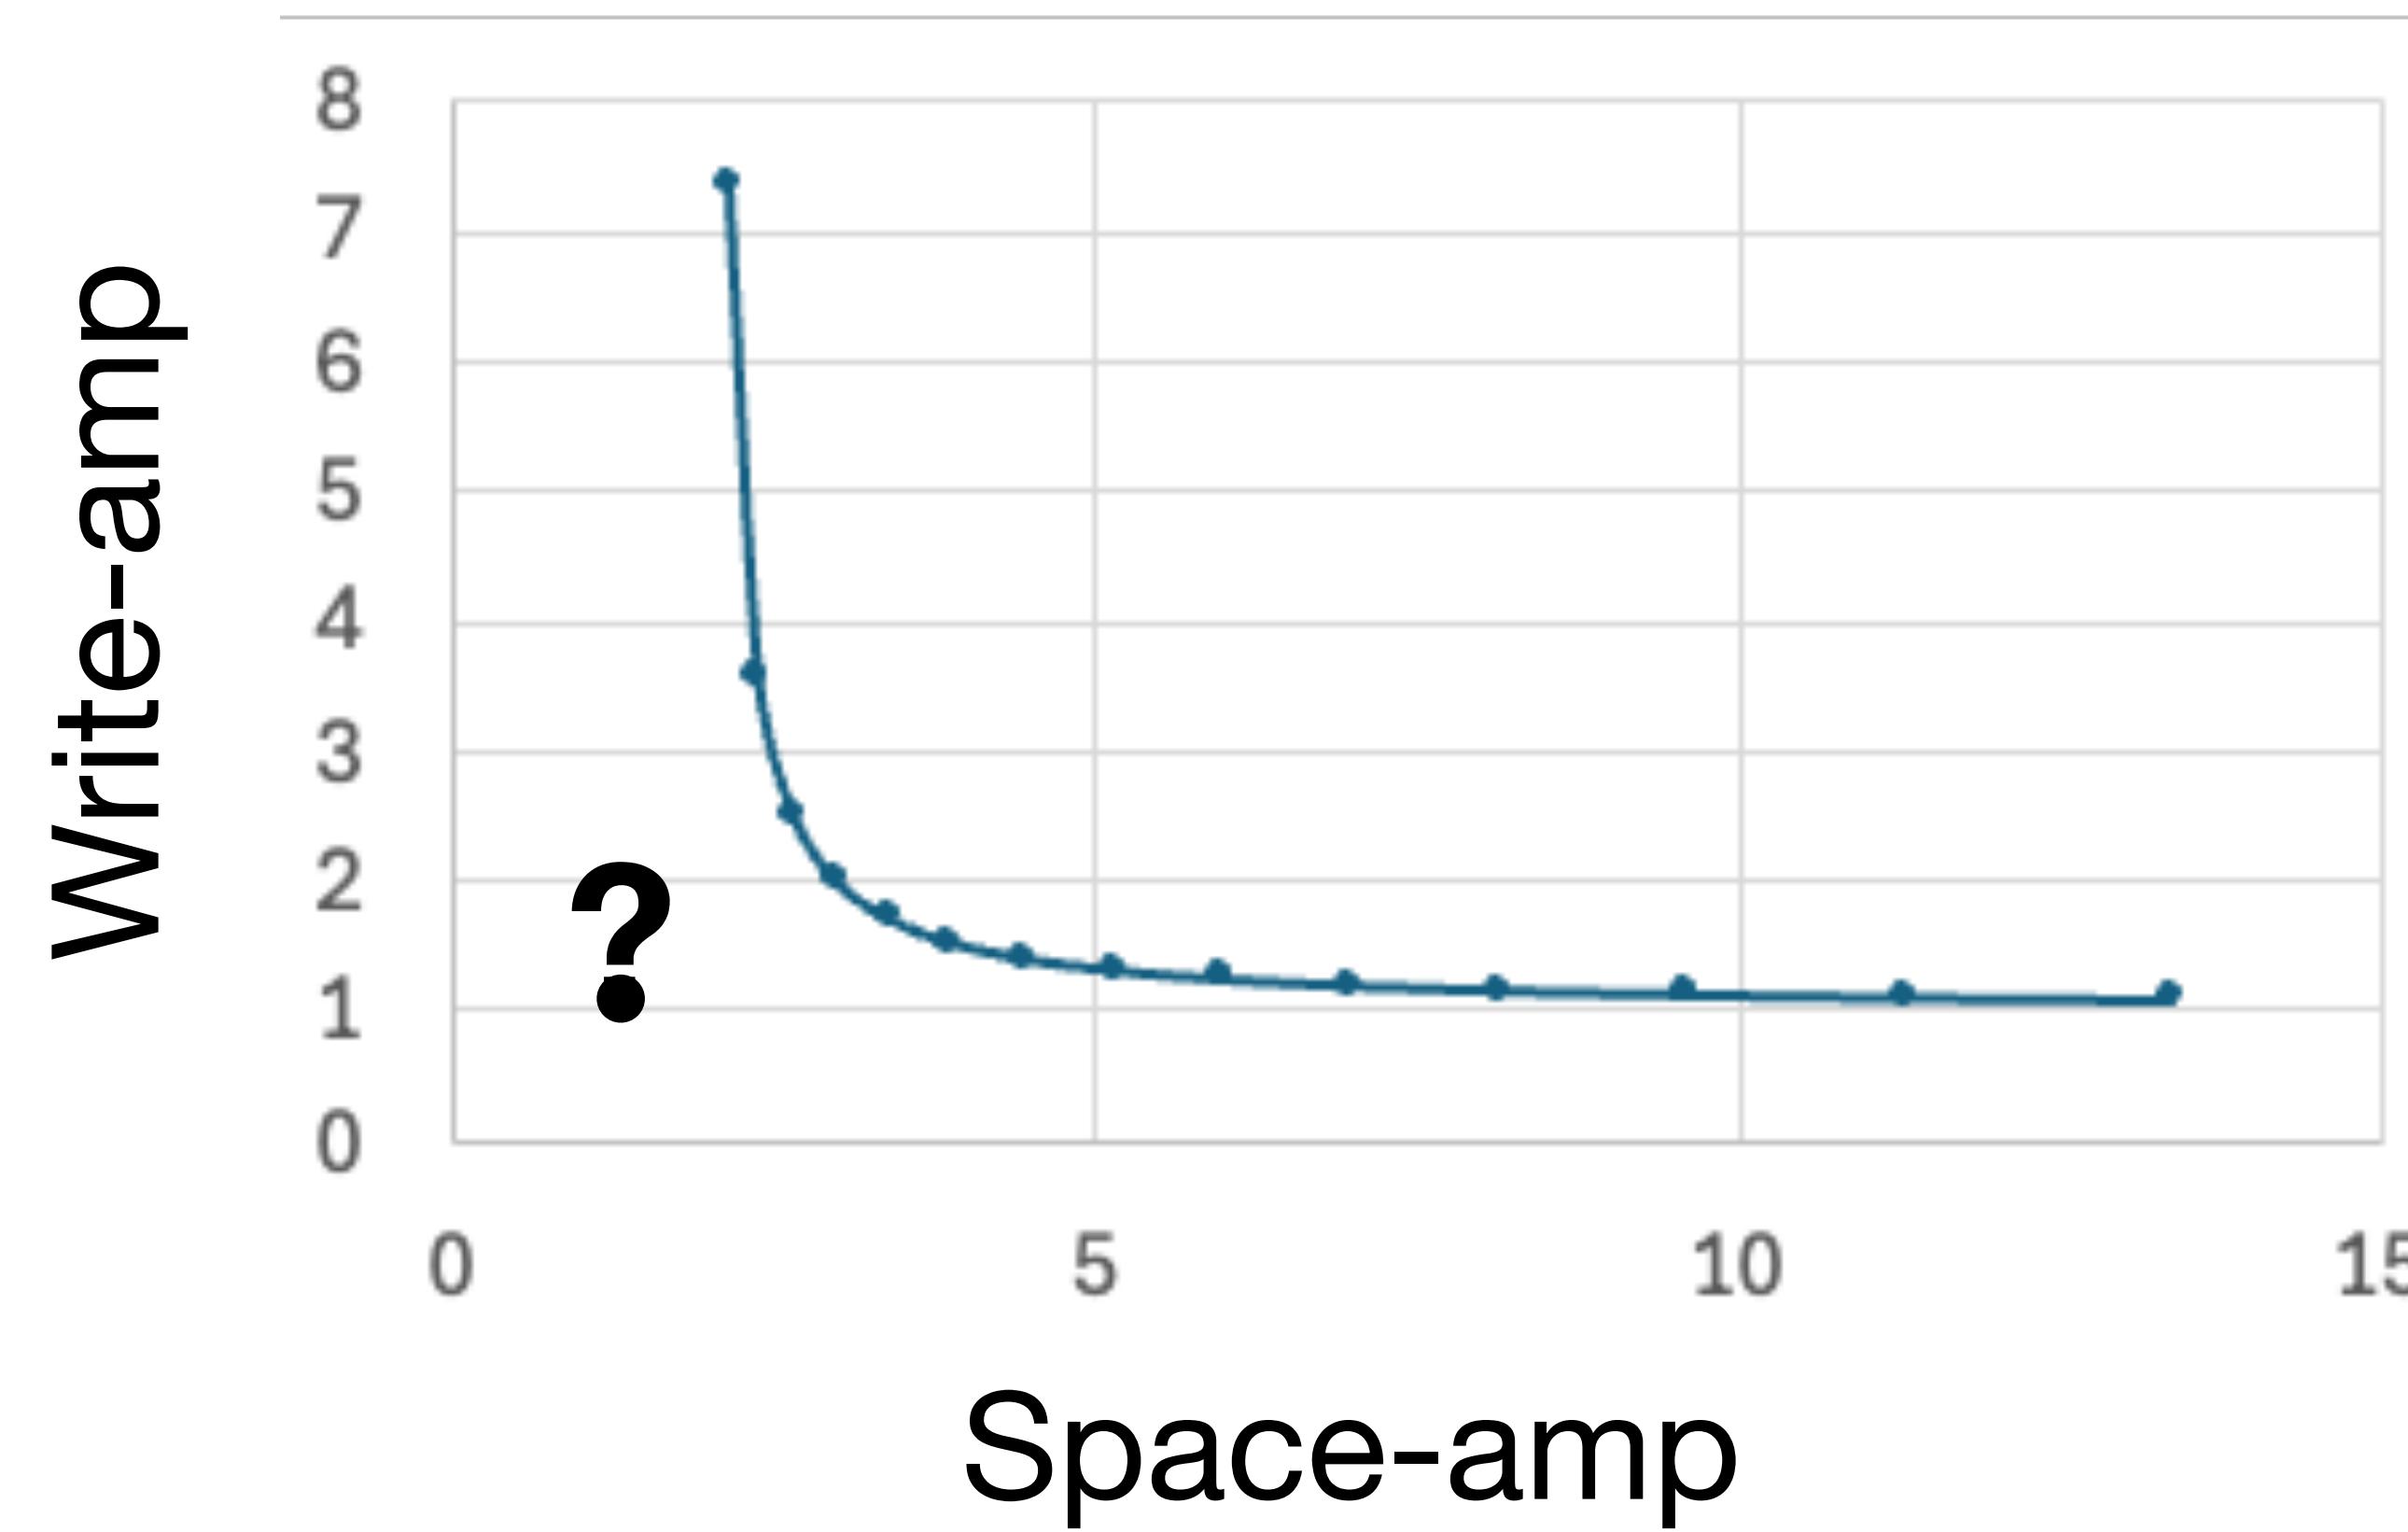
**Reusing Deallocated
Space**



**Alleviating trade-
off via indirection**



Can we completely overcome this trade-off?



Suppose we could expand without copying everything:



Array

Suppose we could expand without copying everything:



Suppose we could expand without copying everything:



Promise:

write-amp of ???

space-amp of ???

Suppose we could expand without copying everything:



Promise:

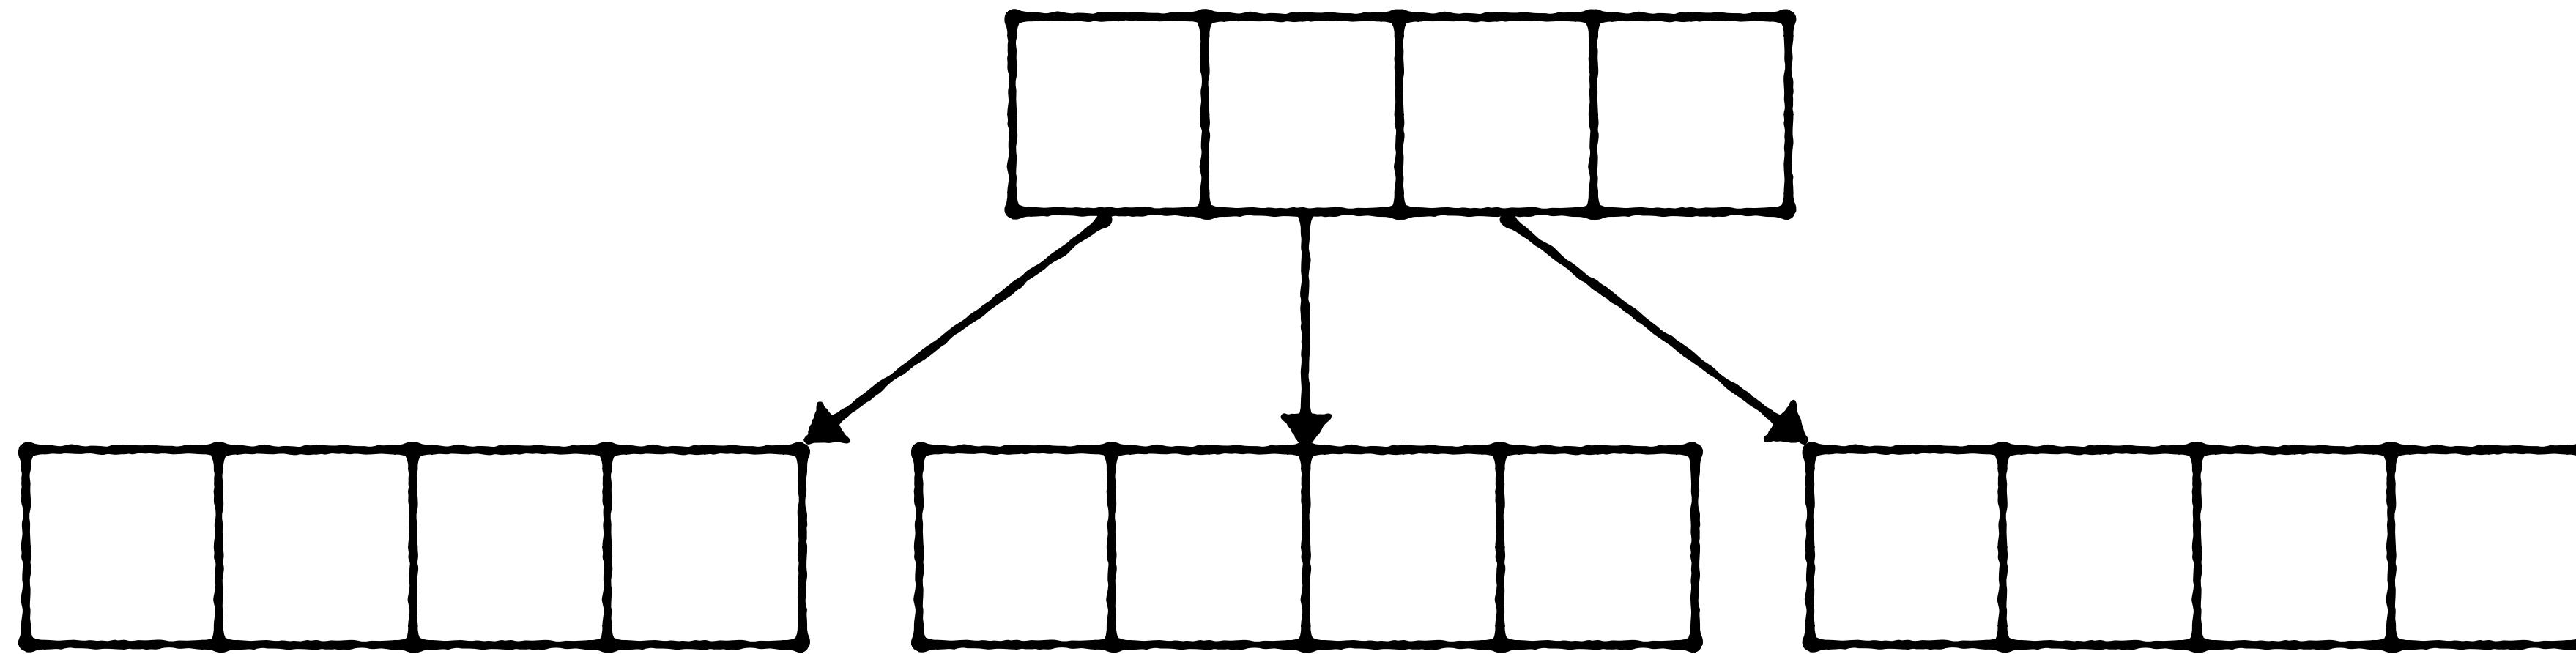
write-amp of ≈ 1

space-amp of ≈ 1

Add a layer of indirection

Directory

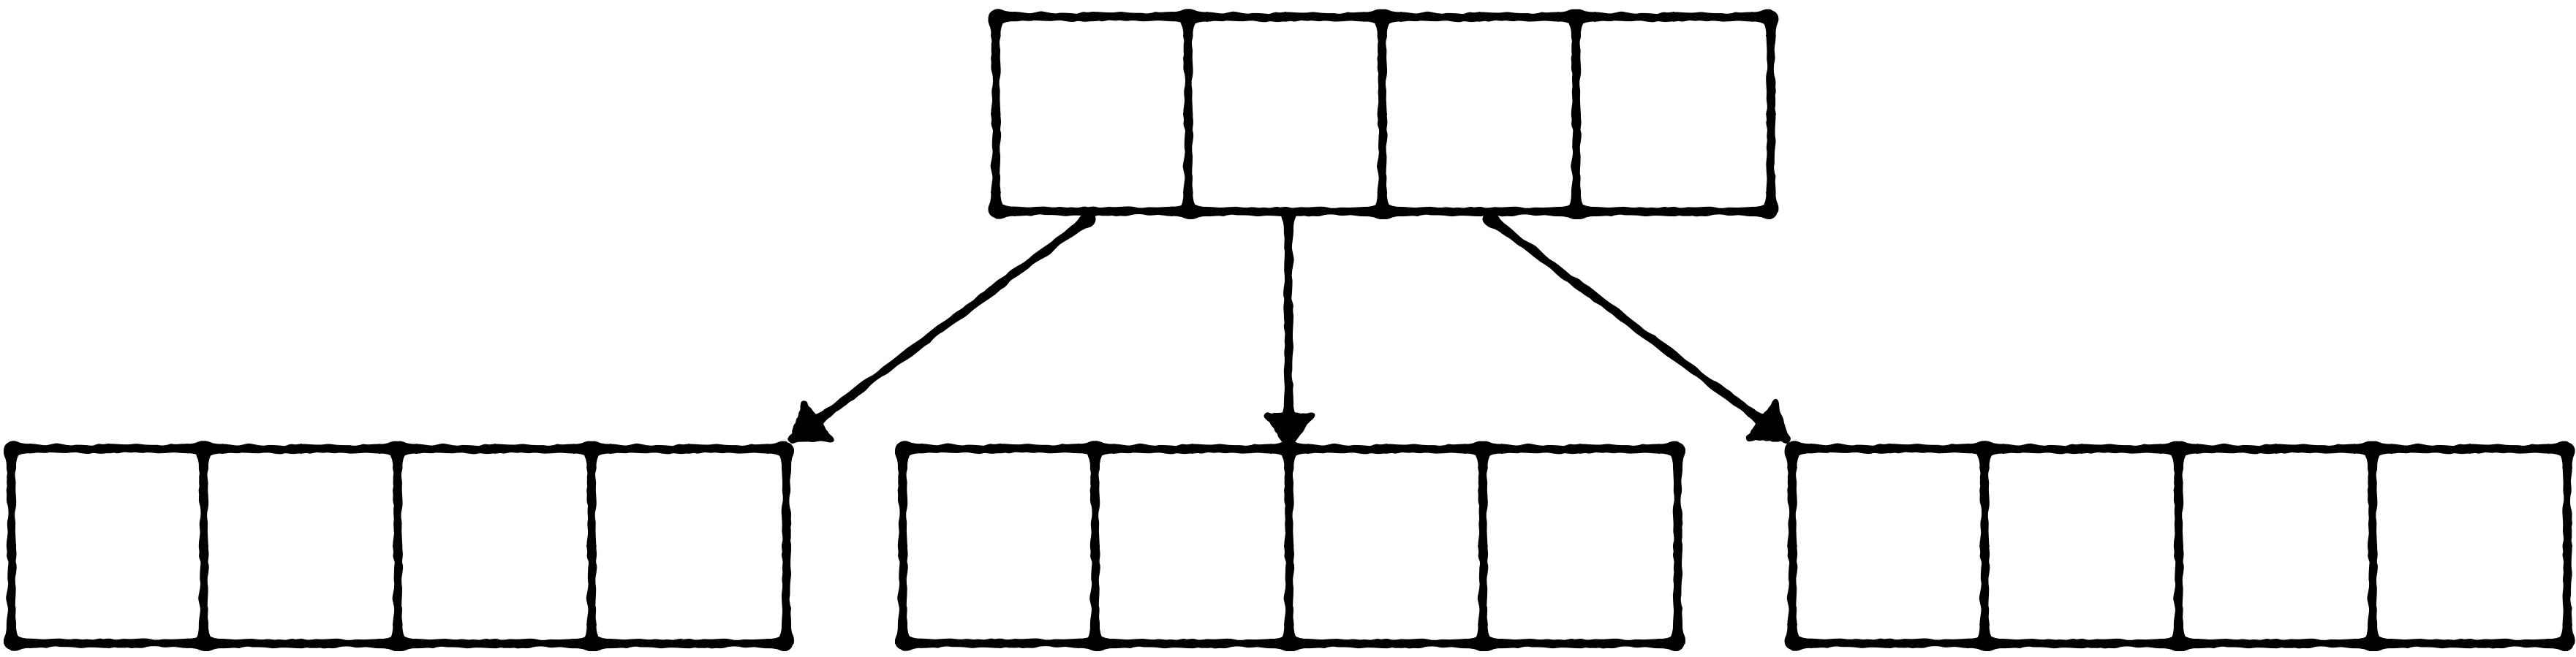
**Data
blocks**



get(i)

Directory

Data
blocks

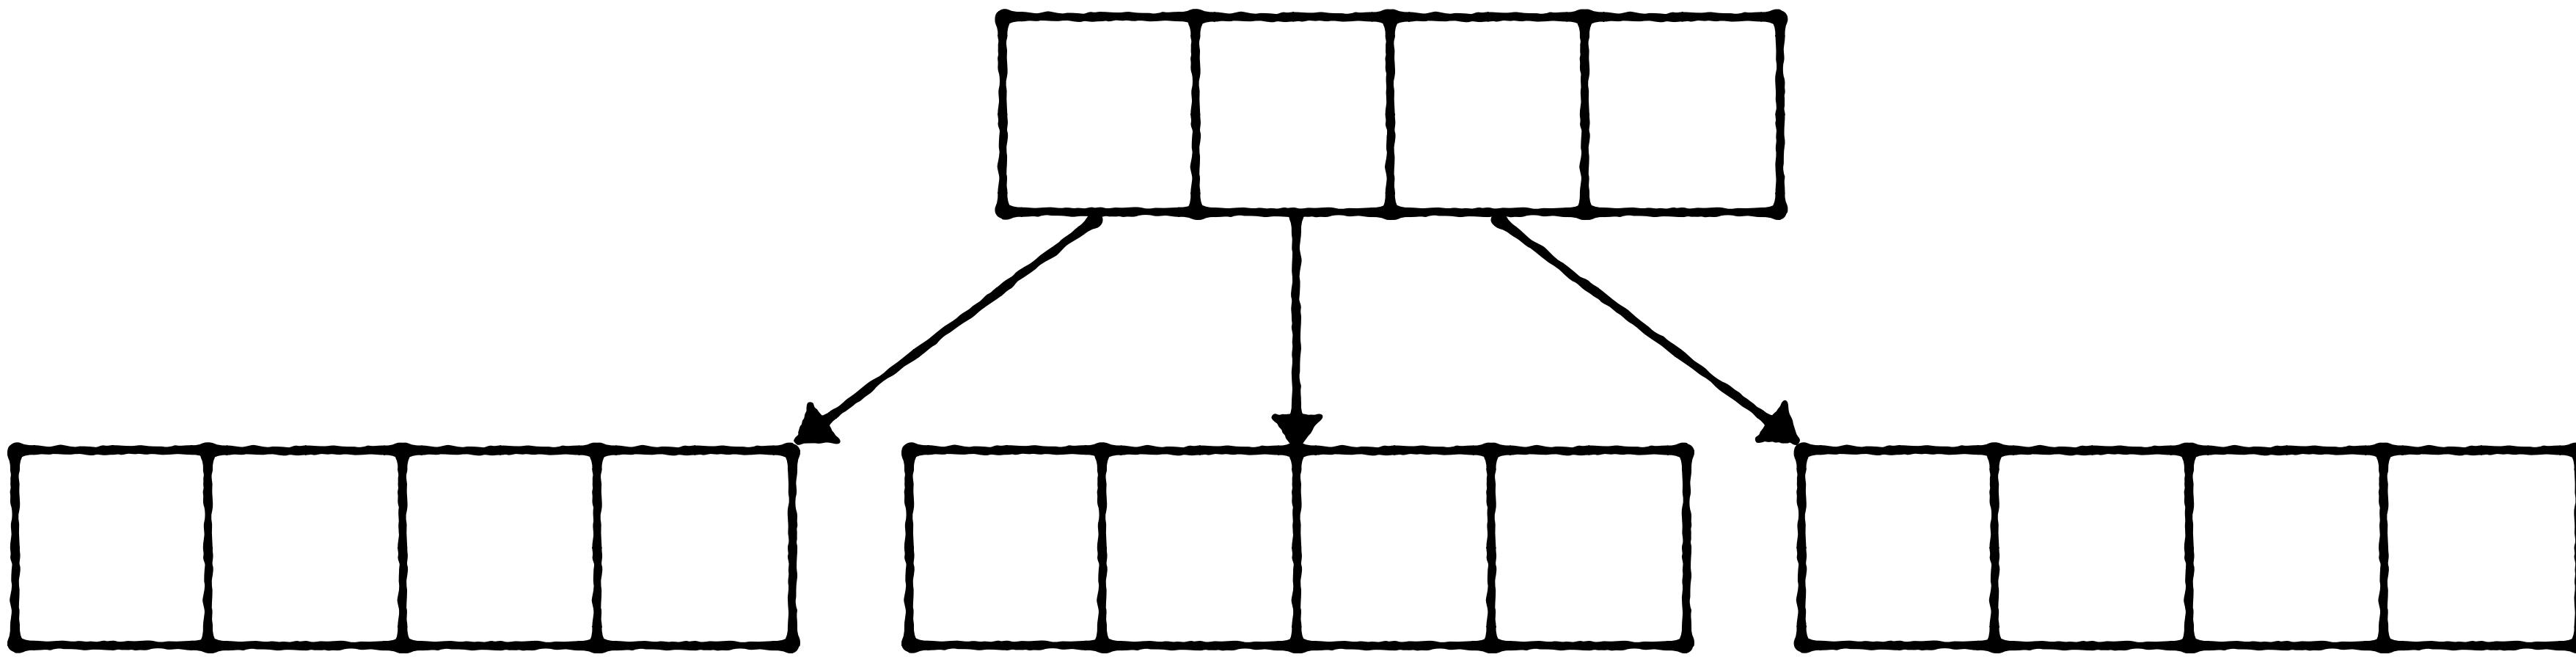


get(i)

Data block = $\lfloor i / \text{data block size} \rfloor$

Directory

Data
blocks

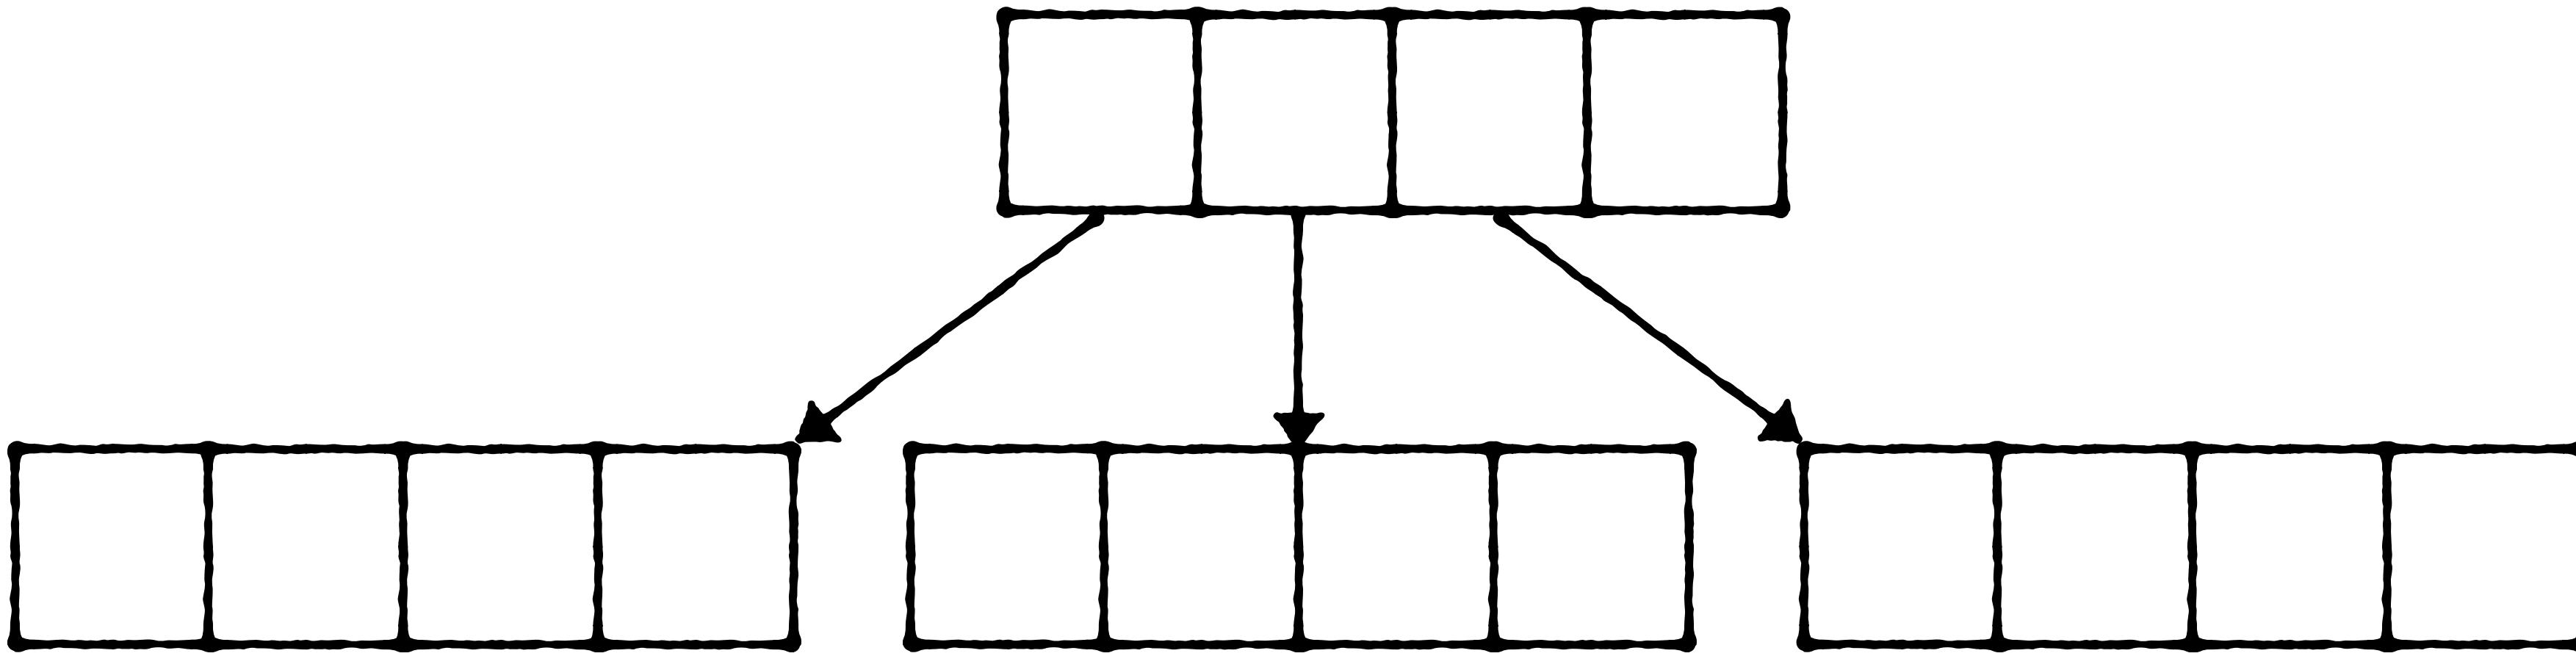


get(i)

Data block = $\lfloor i / \text{data block size} \rfloor$
offset within = $i \% \text{ data block size}$

Directory

Data
blocks

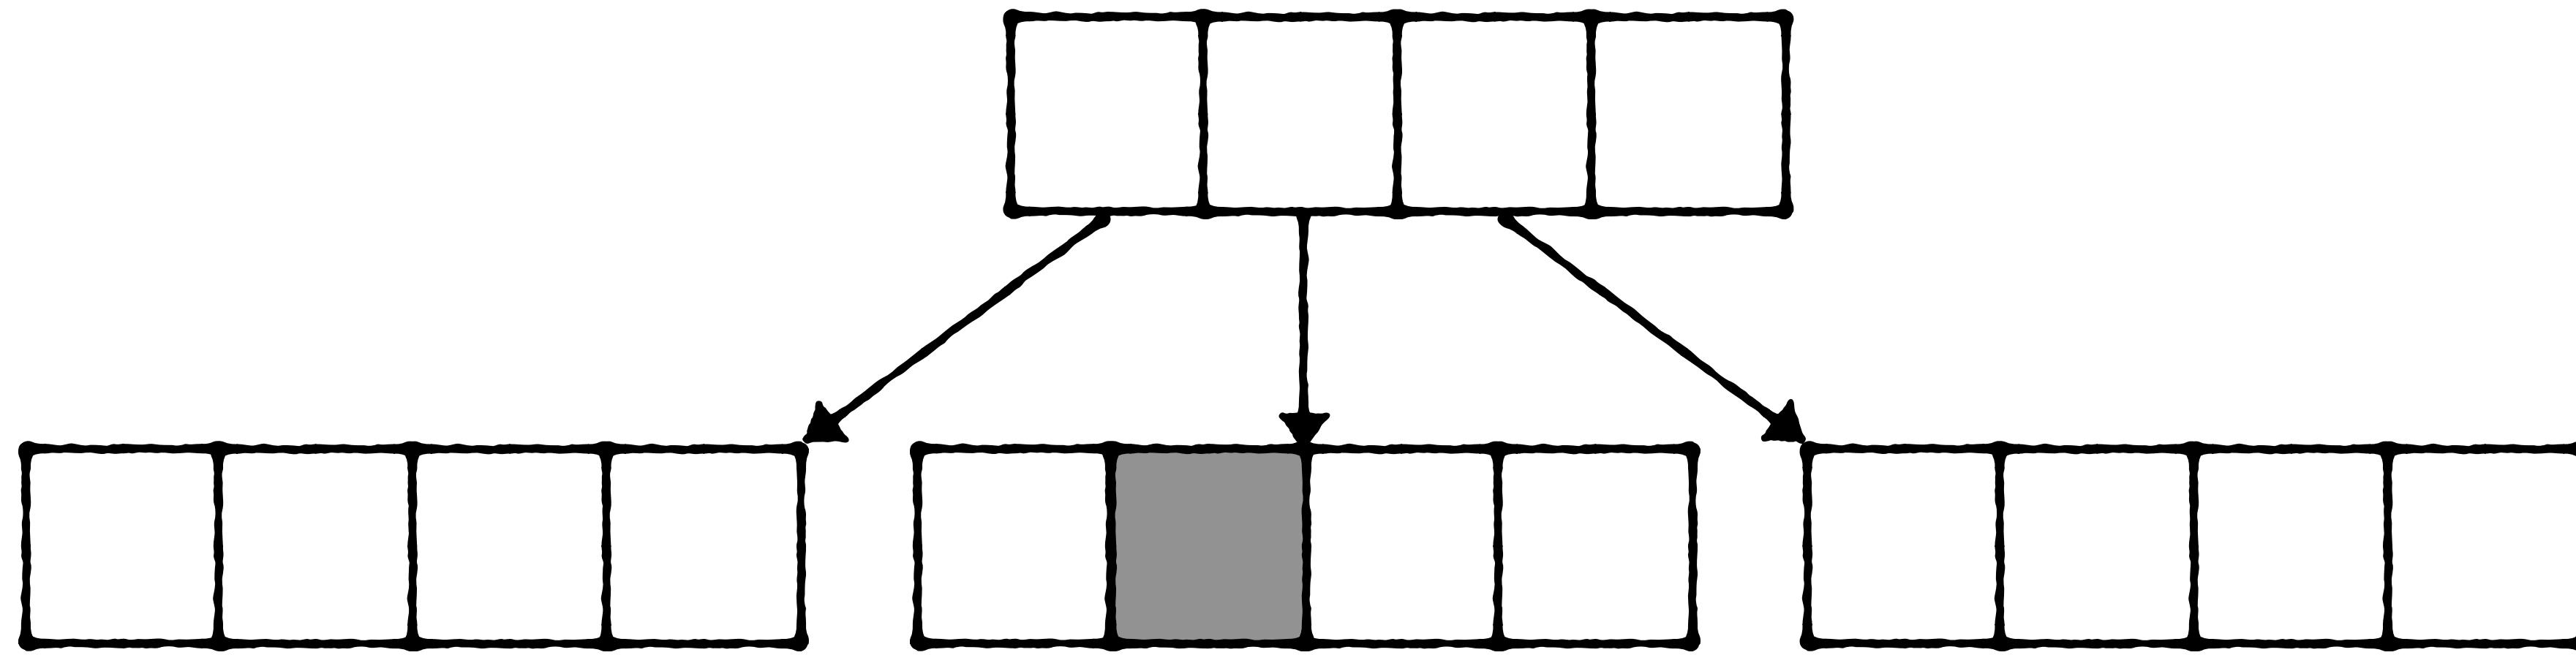


get(5)

Data block = $\lfloor 5 / 4 \rfloor = 1$
offset within = $5 \% 4 = 1$

Directory

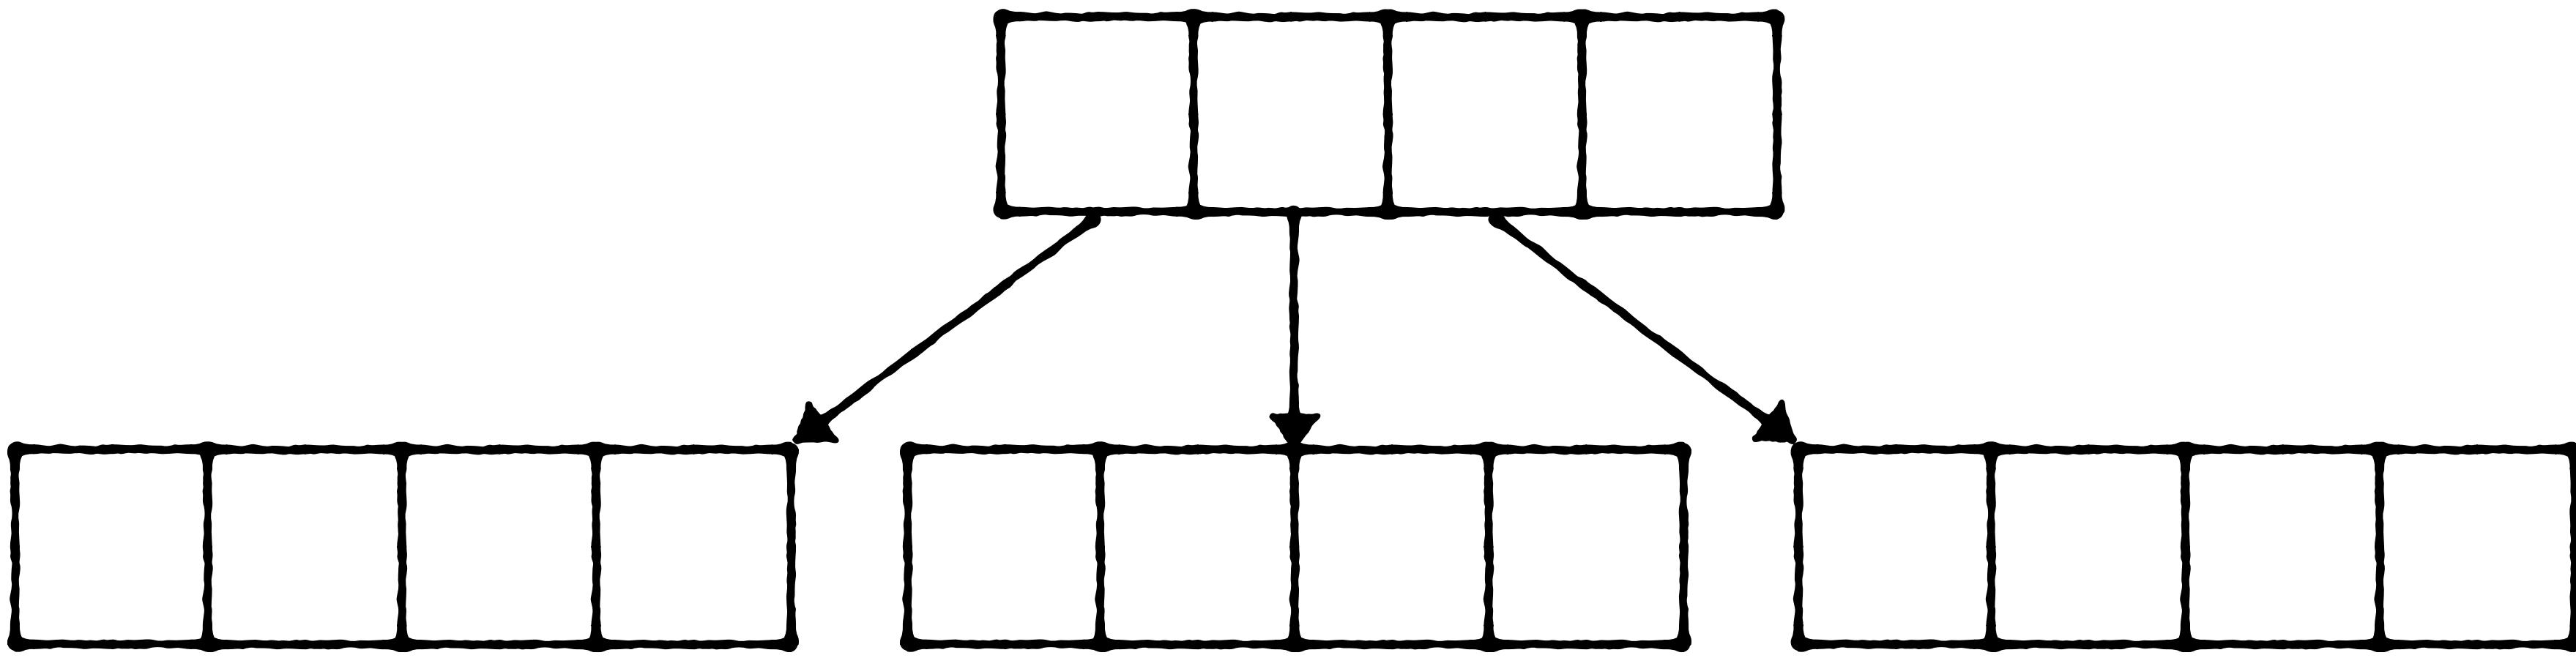
Data
blocks



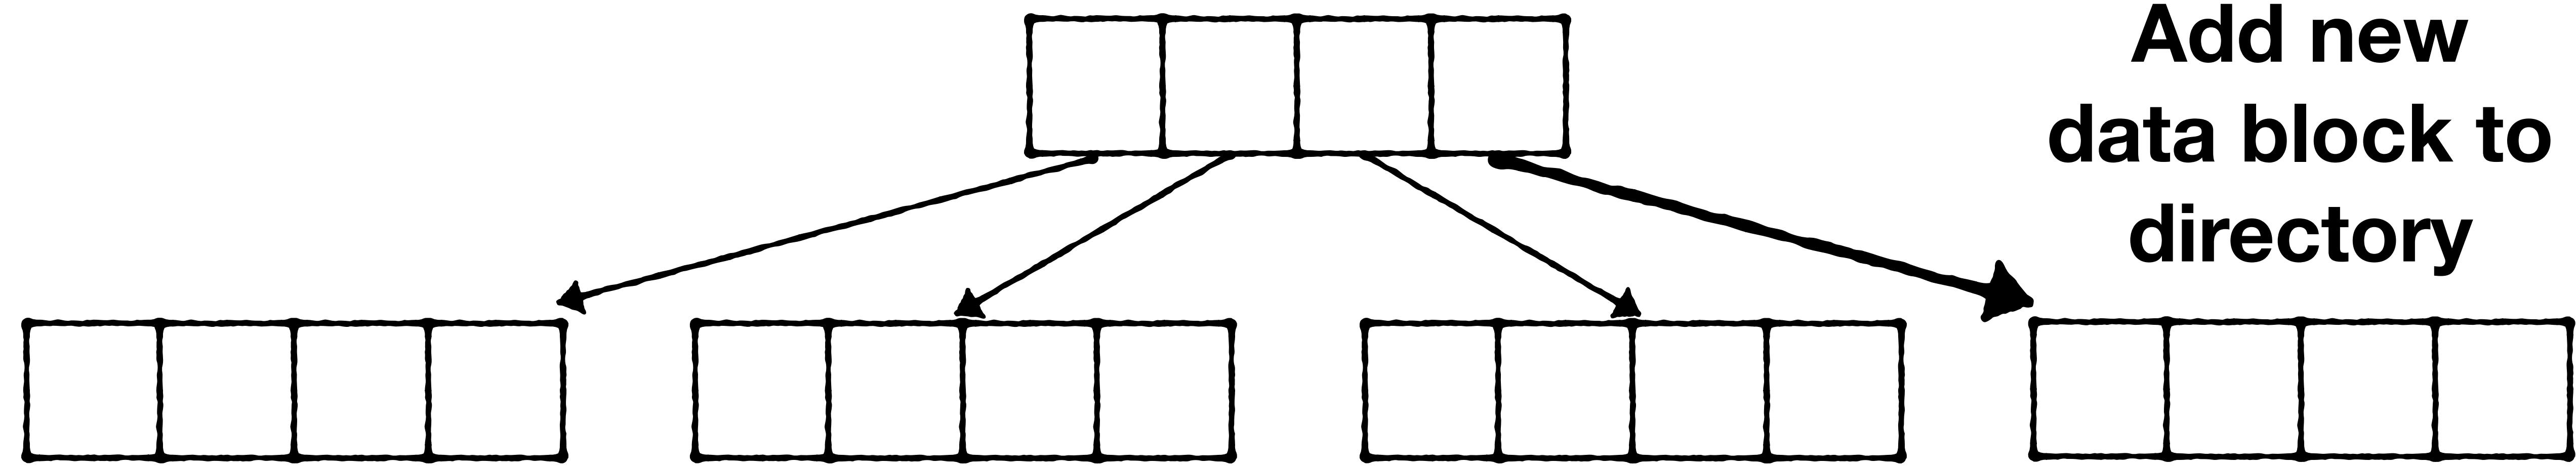
Expand?

Directory

Data
blocks

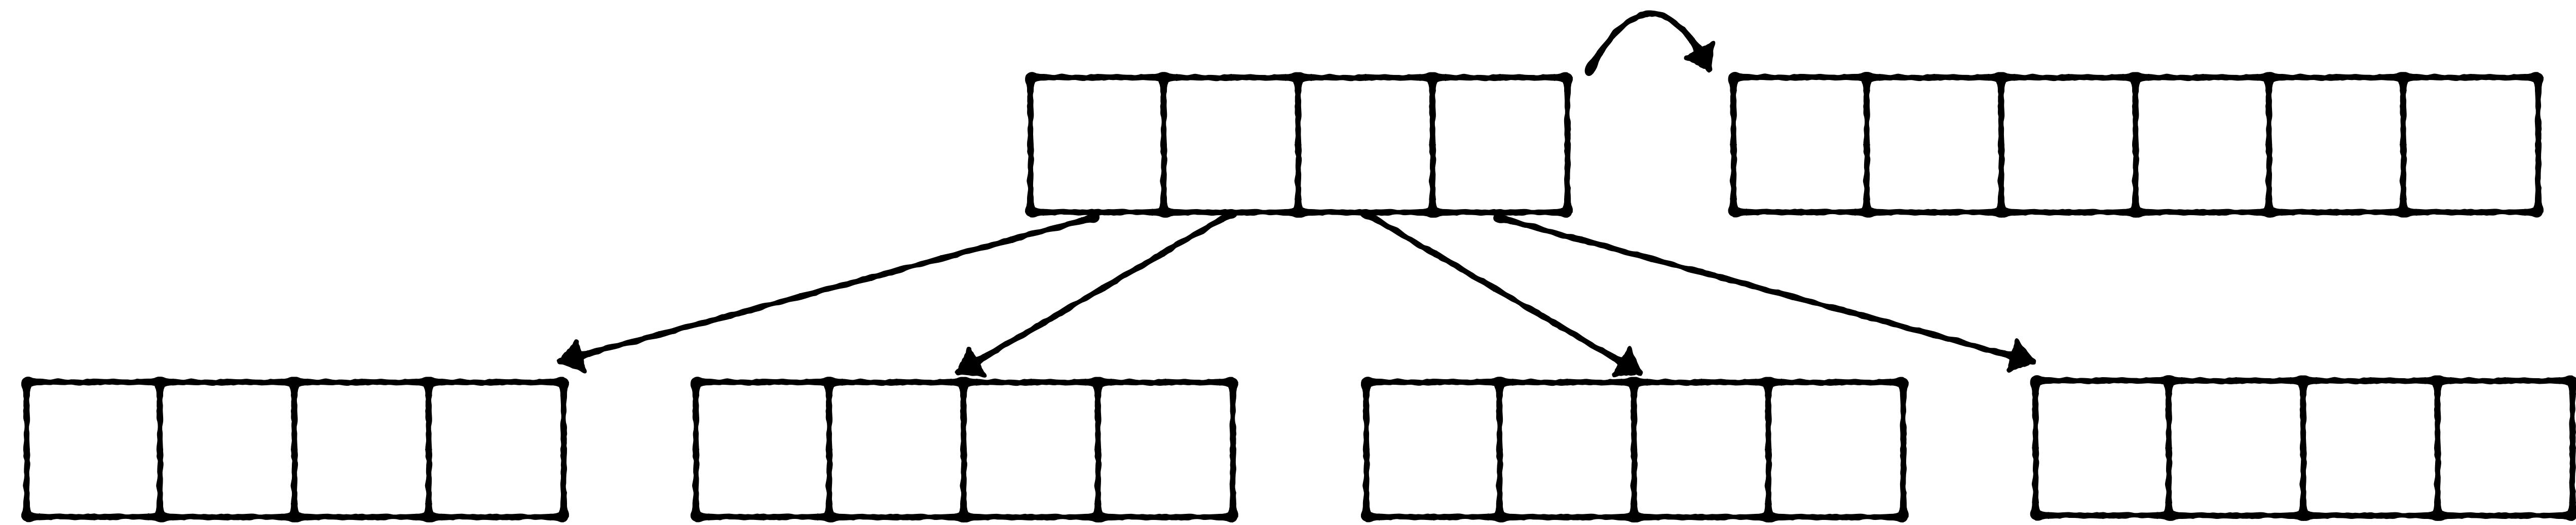


Expand?



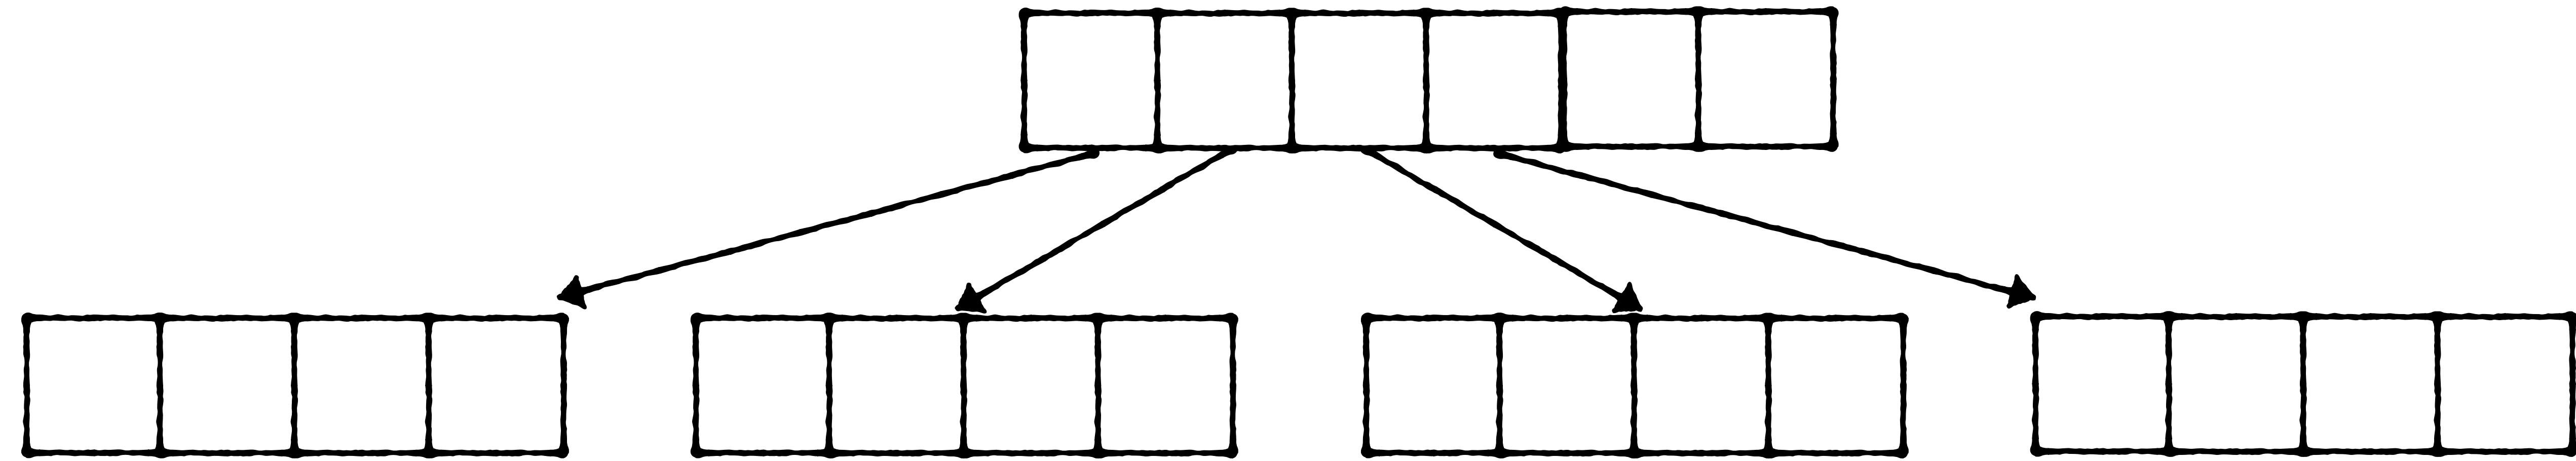
Expand?

**Expand directory if
we need more space**

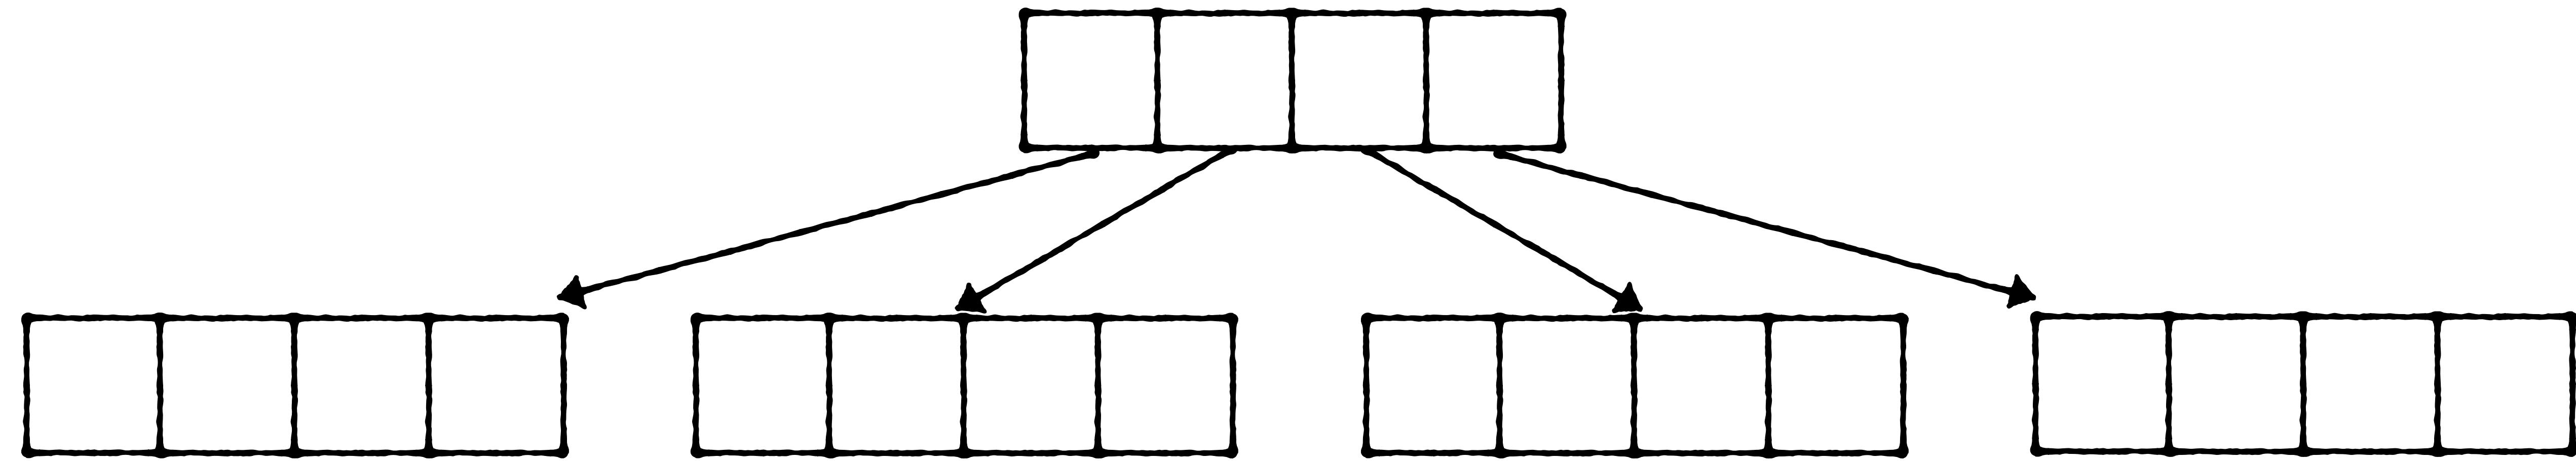


Expand?

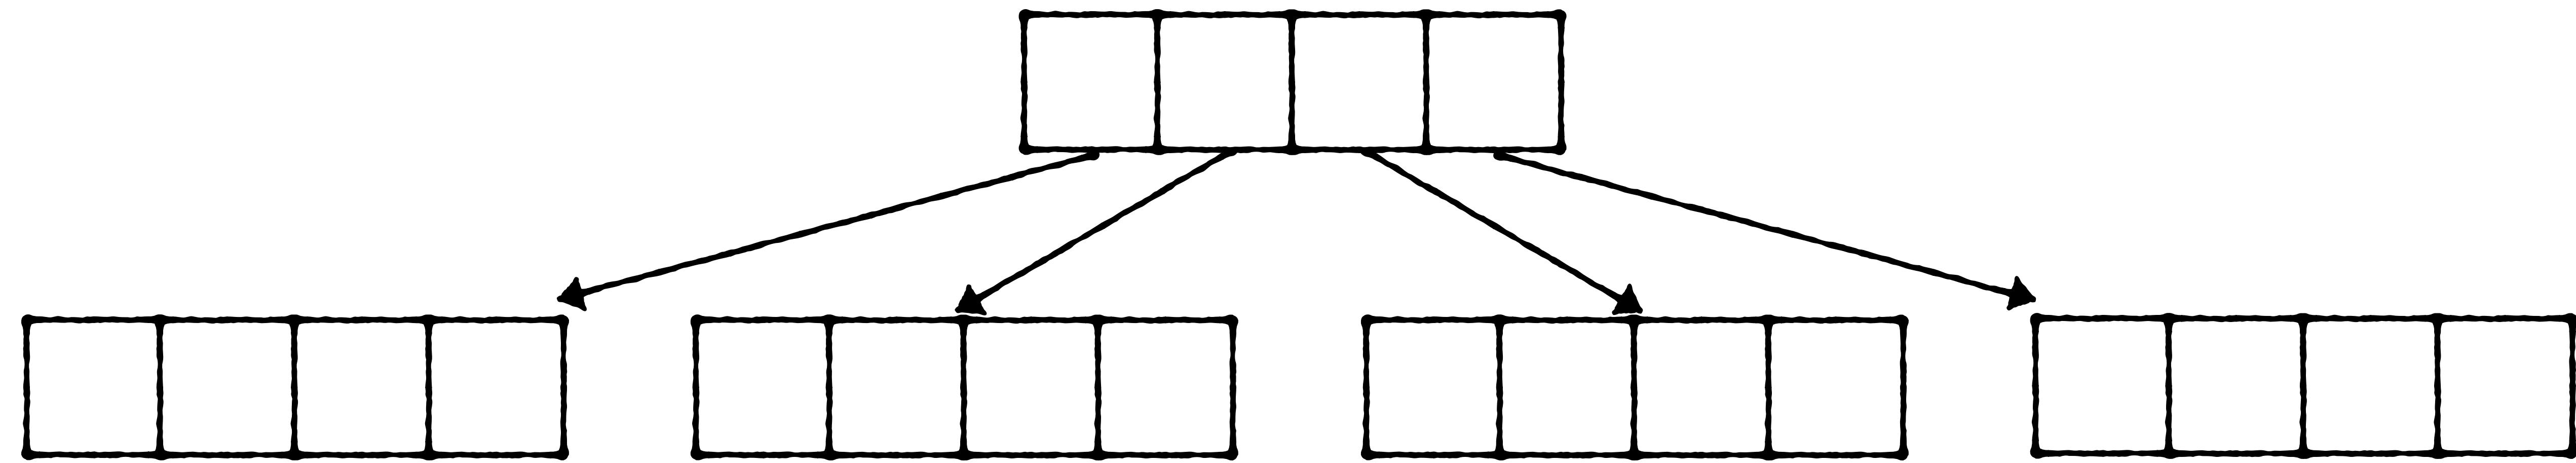
**Expand directory if
we need more space**



Downside?

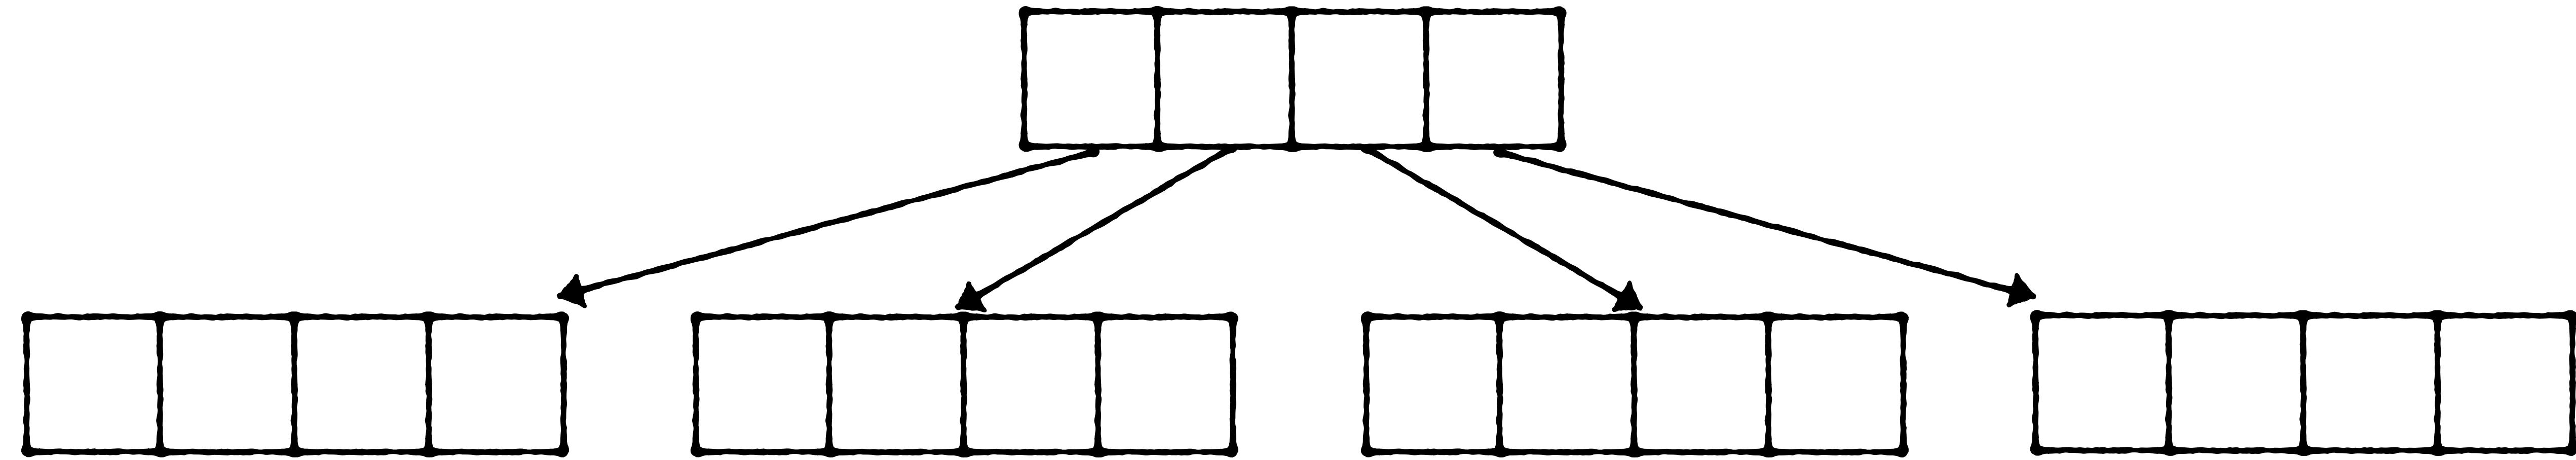


Downside: 2 memory hops per access



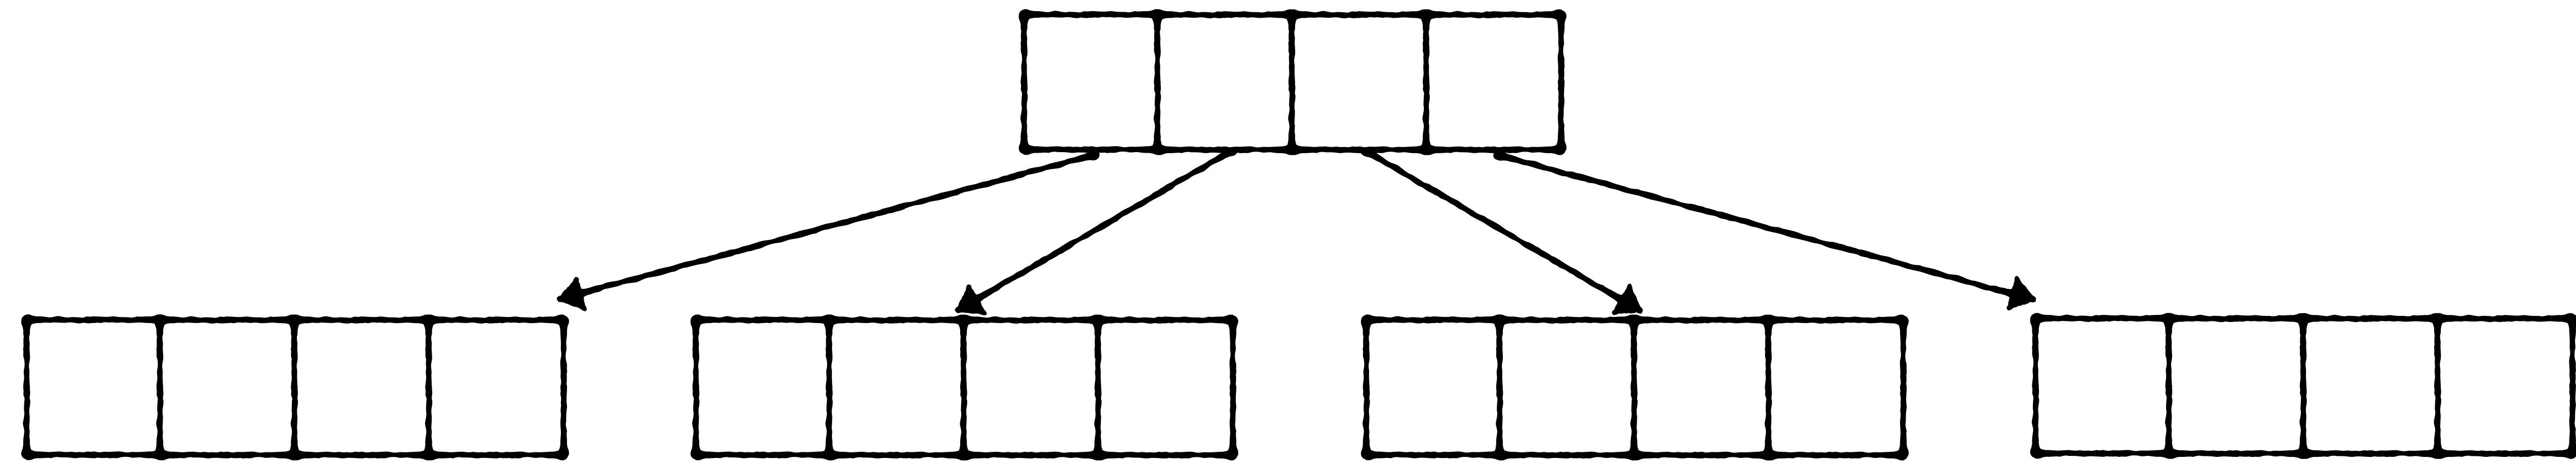
Downside: 2 memory hops per access

Mitigation?



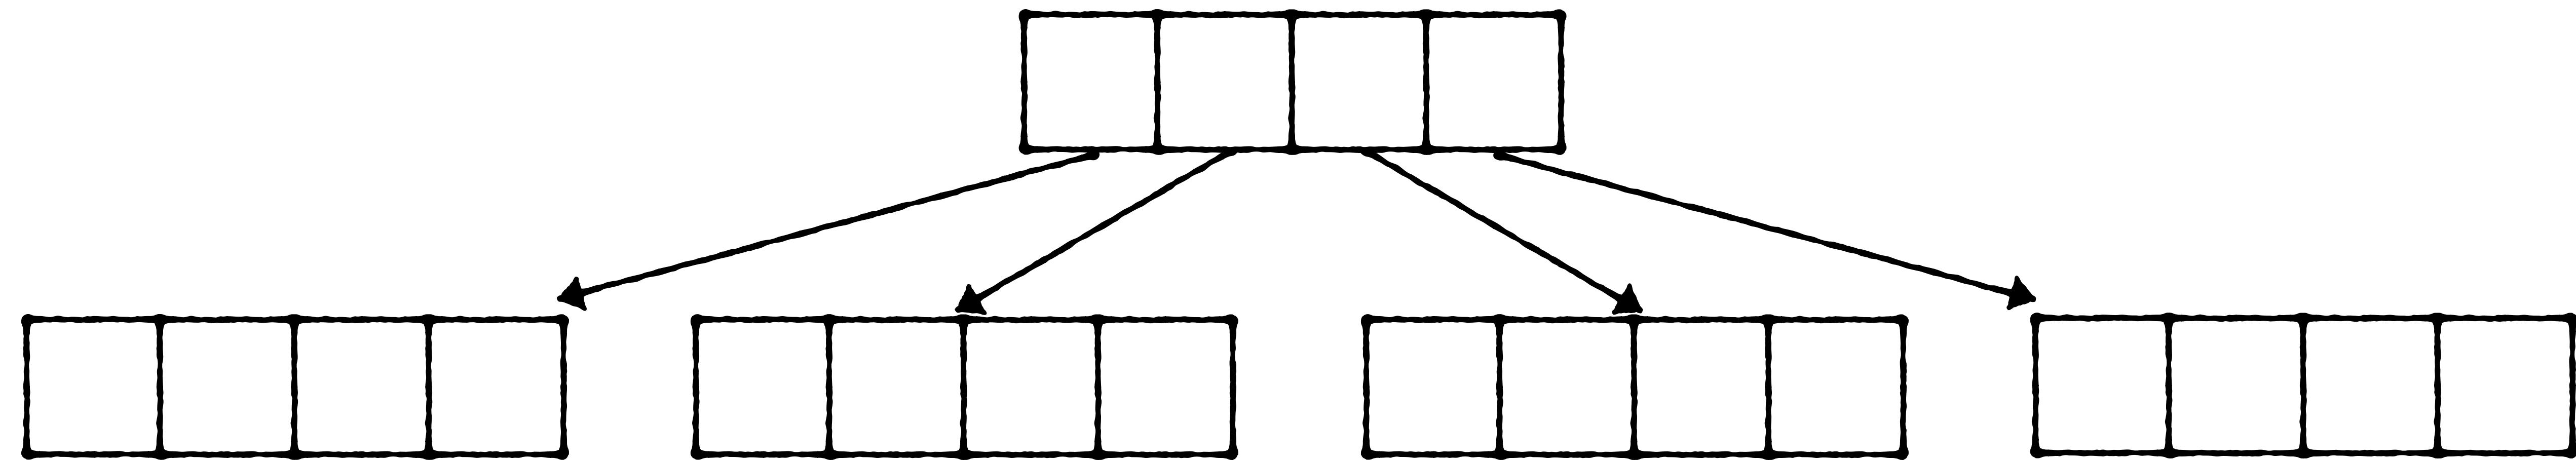
Downside: 2 memory hops per access

Mitigation: directory must fit in L1 cache



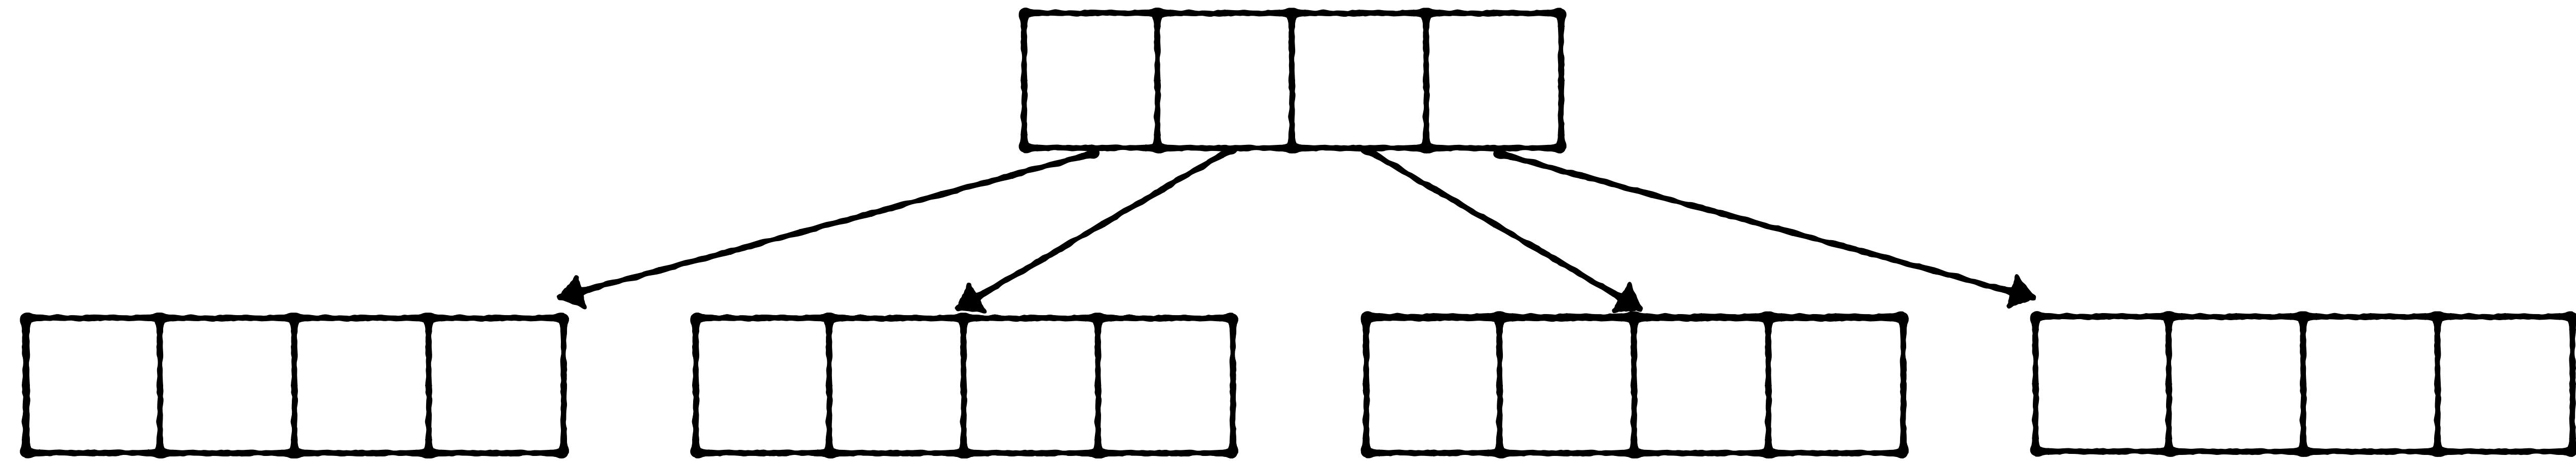
Downside: 2 memory hops per access

Mitigation: directory must fit in L1 cache

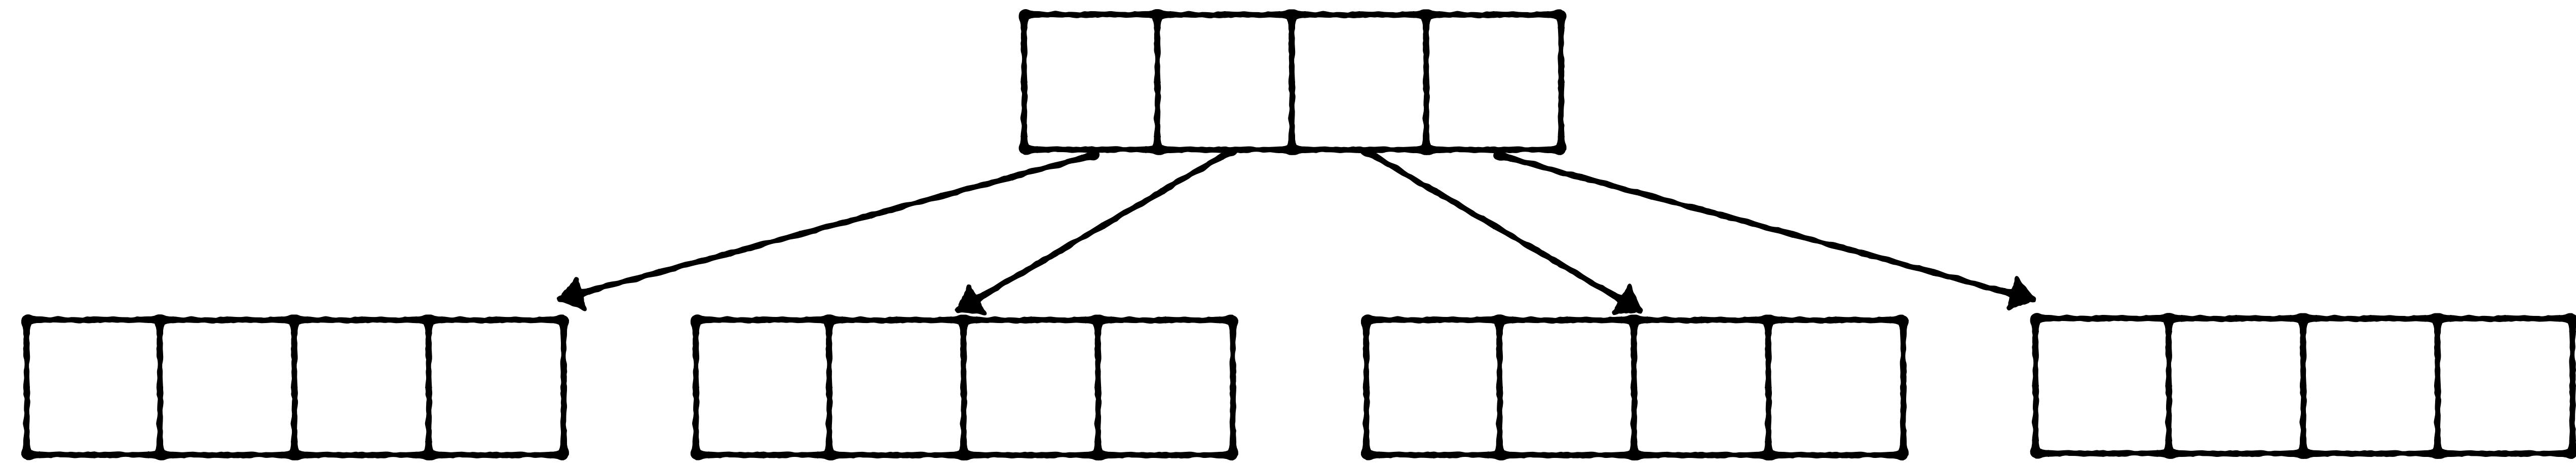


**Typical L1 cache size:
16-128 KB per core**

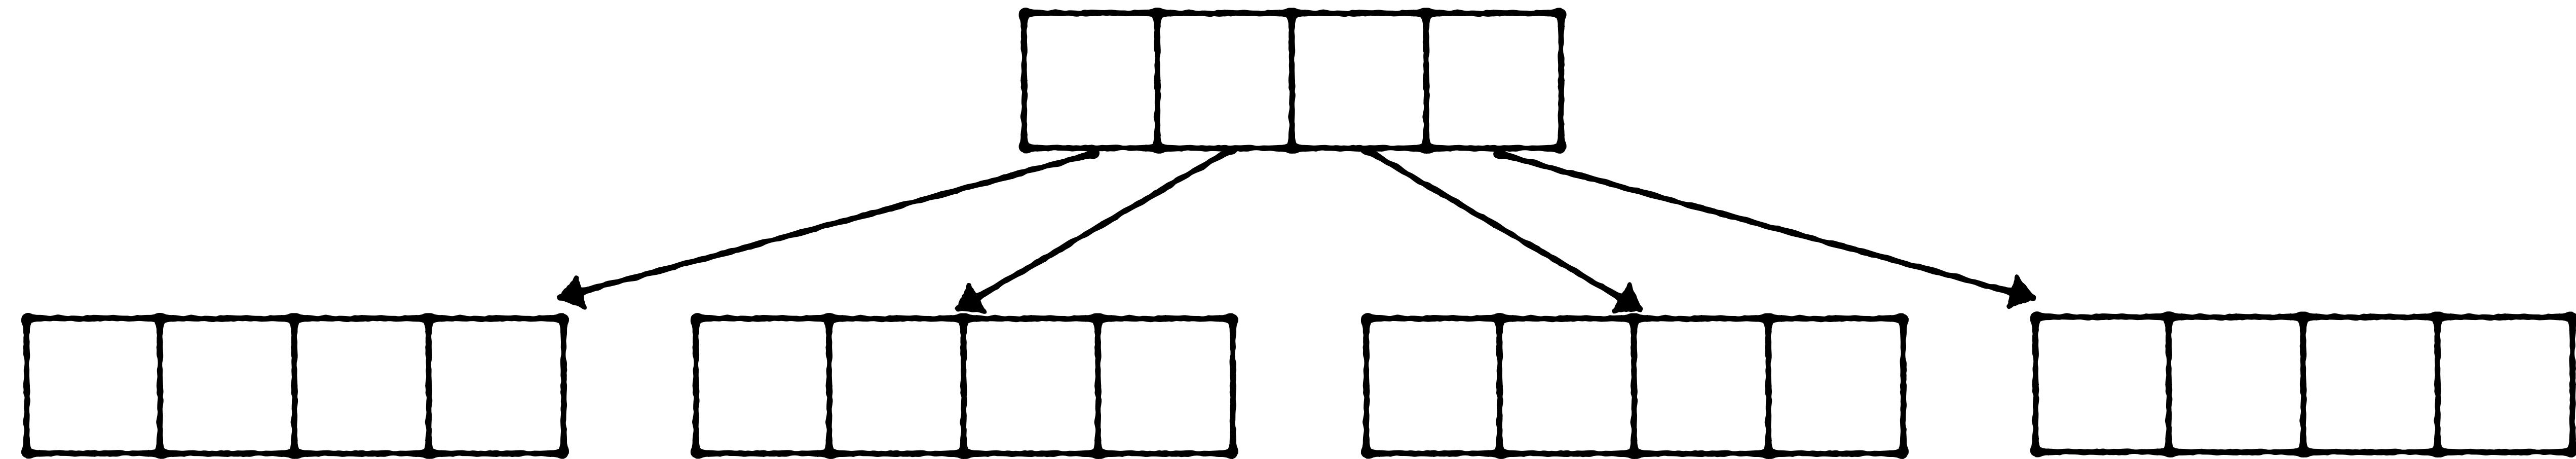
directory size?



$$\text{directory size} = \frac{\text{Data size}}{\text{Data block size}}$$

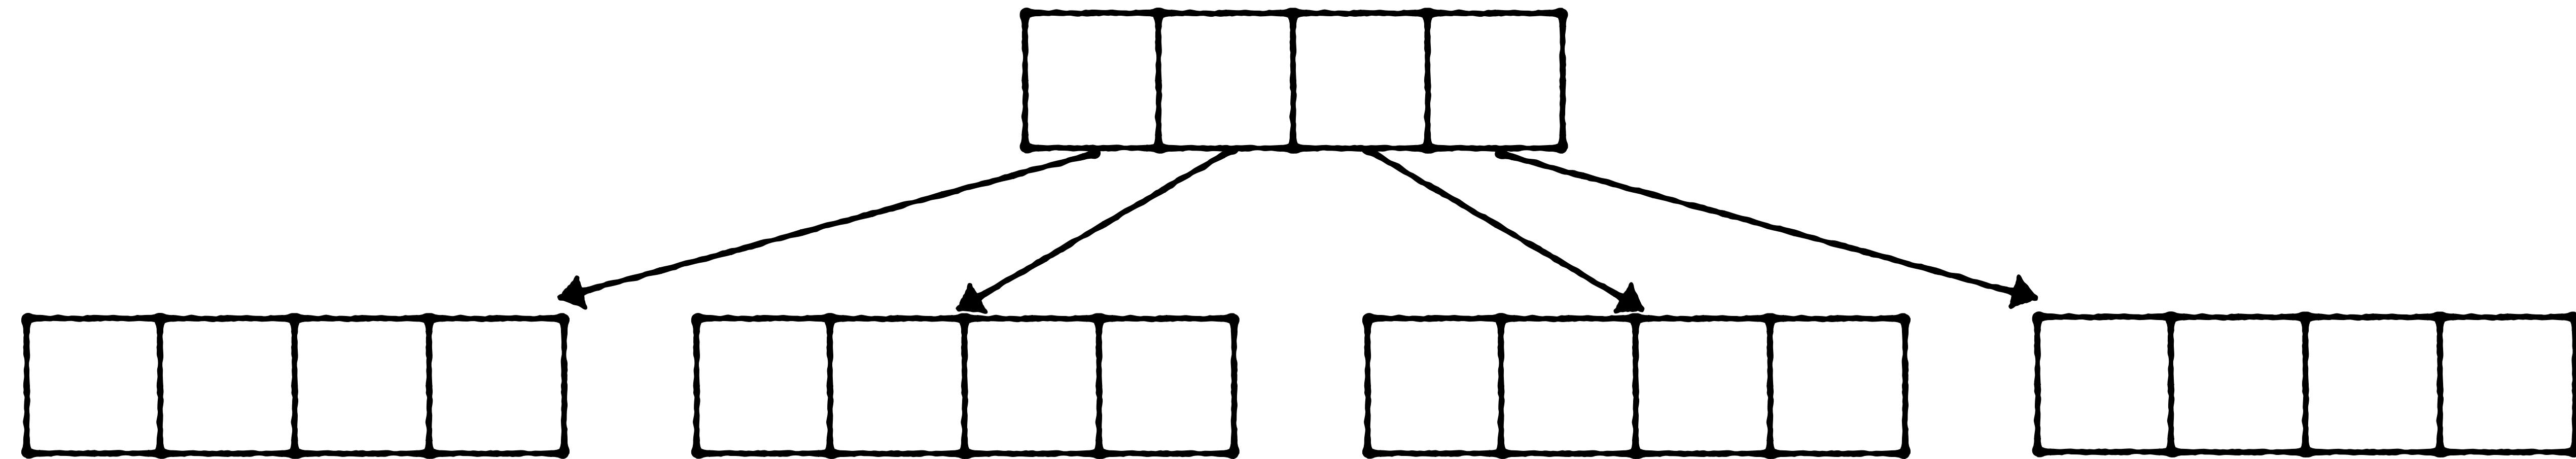


$$\text{directory size} = \frac{\text{Data size}}{\text{Data block size}} = O(N)$$



Risk: data blocks are initialized too small

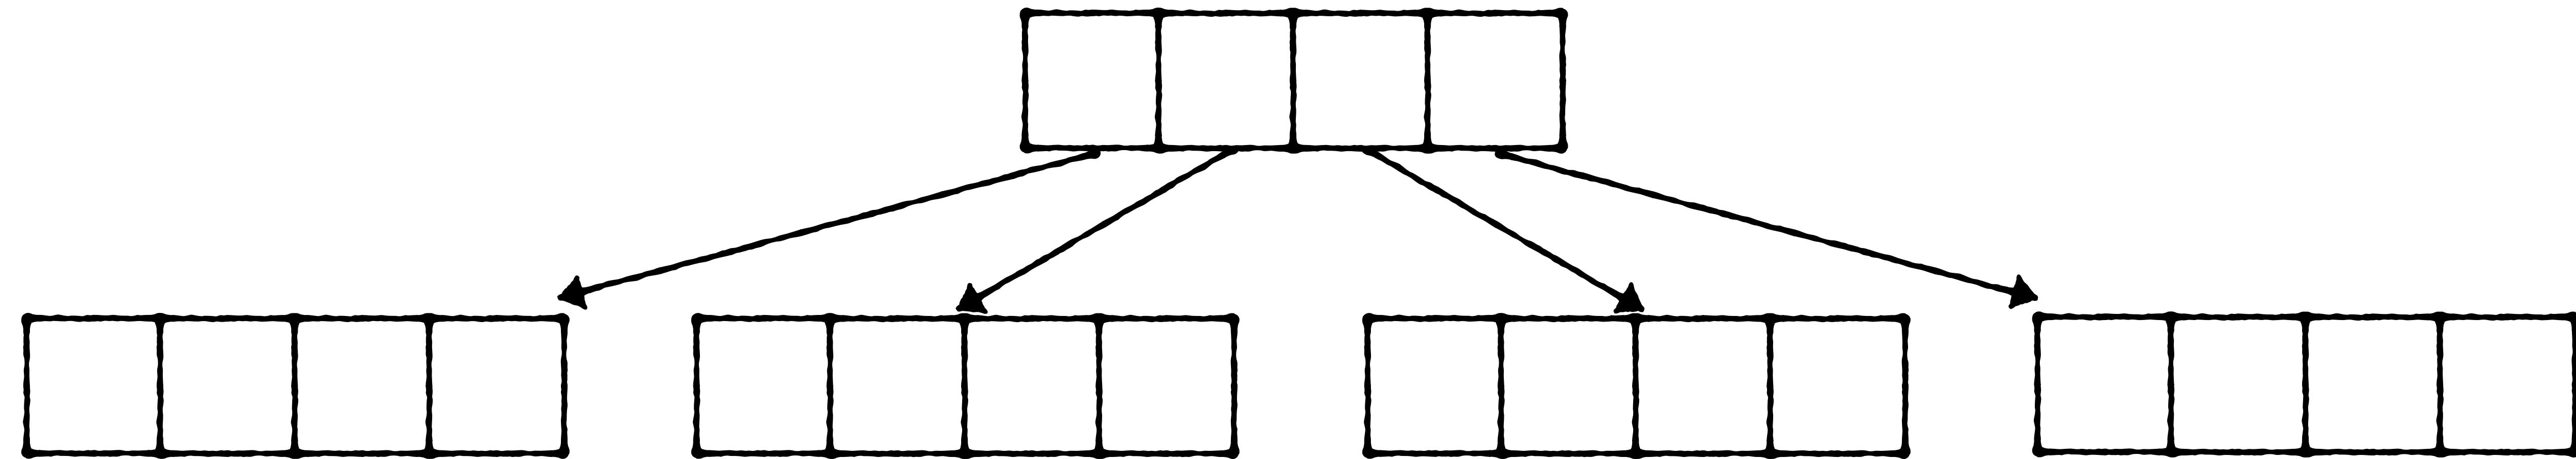
$$\text{directory size} = \frac{\text{Data size}}{\text{Data block size}} = O(N)$$



Risk: data blocks are initialized too small

Directory may outgrow the L1 cache

$$\text{directory size} = \frac{\text{Data size}}{\text{Data block size}} = O(N)$$



Risk: data blocks are initialized too small

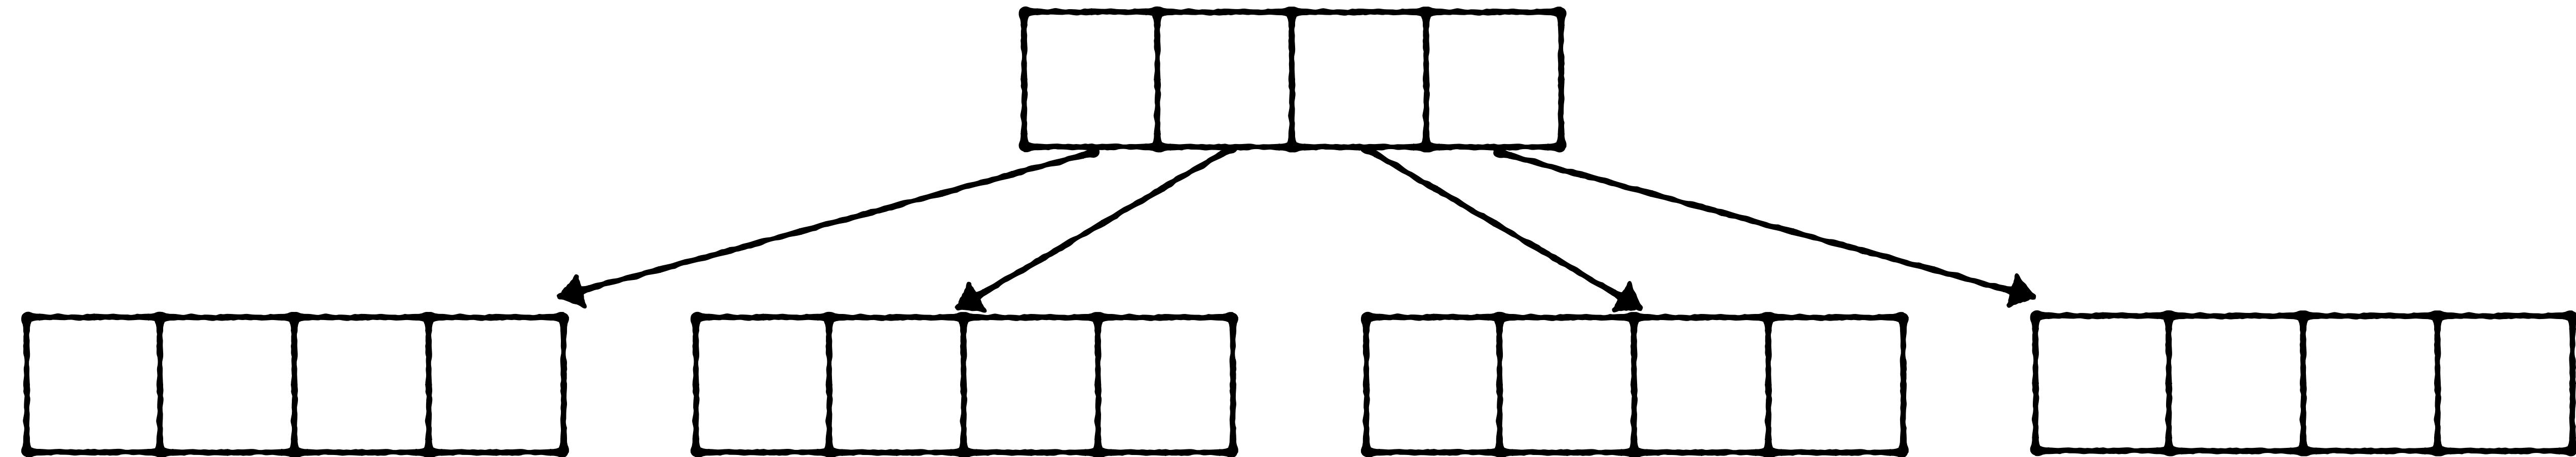
Directory may outgrow the L1 cache

Solution?

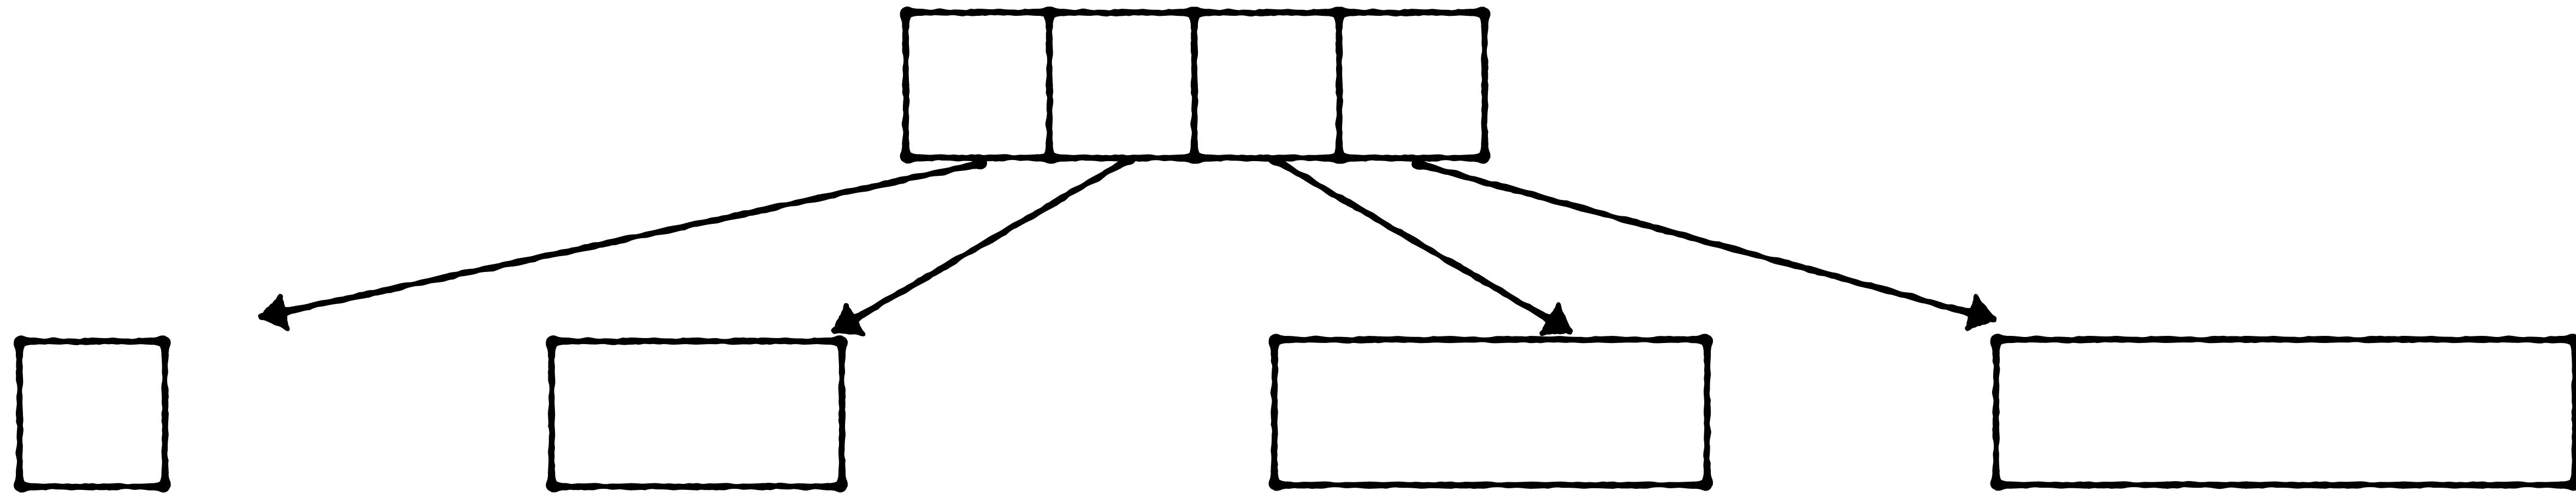
Resizable Arrays in Optimal Time and Space

Algorithms and Data Structures Symposium, 1999

Andrej Brodnik, Svante Carlsson, Erik D. Demaine, J. Ian Munro, and Robert Sedgewick

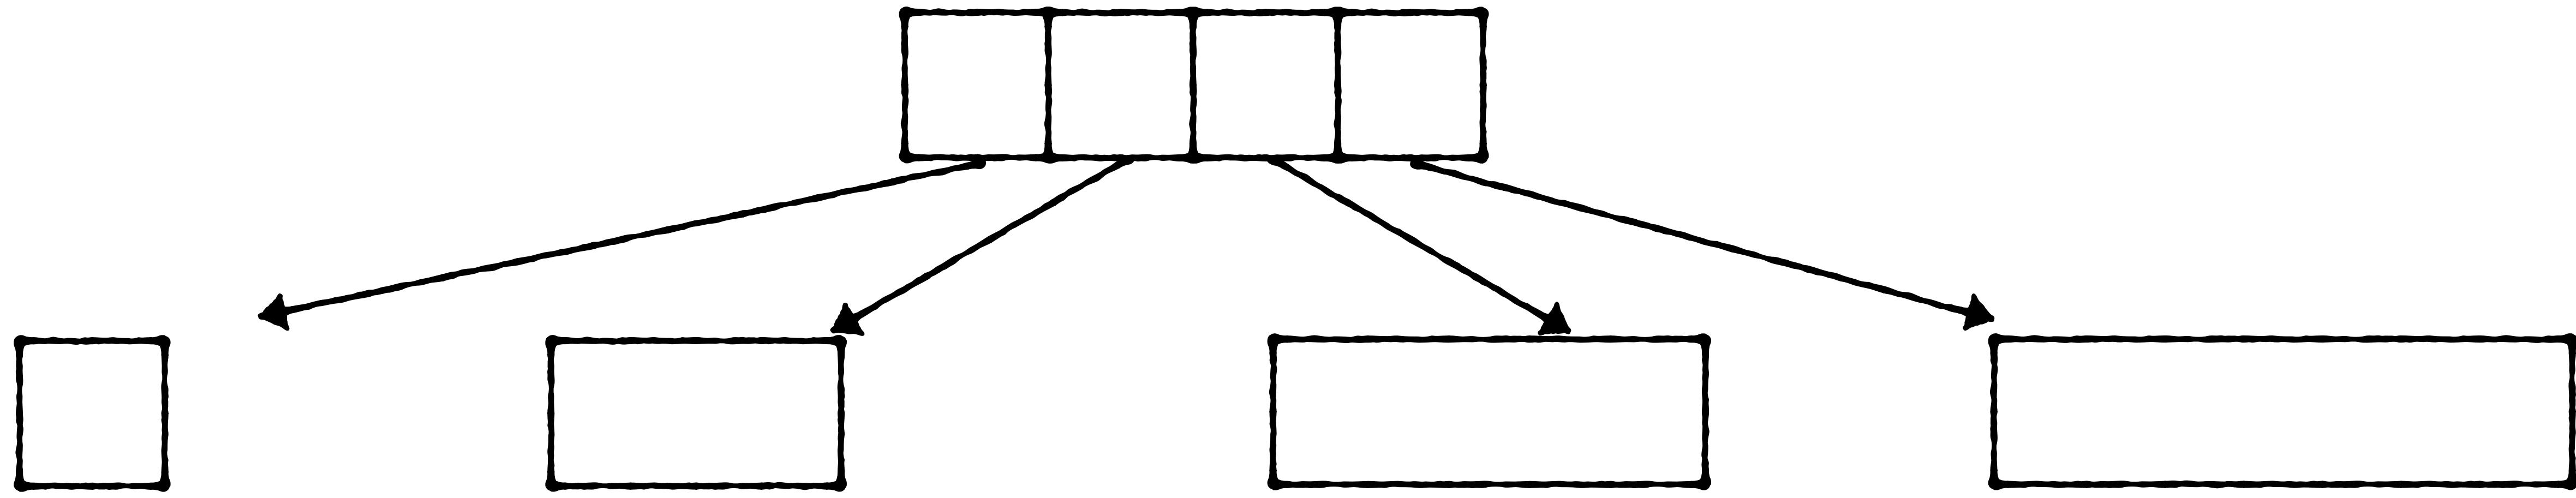


Resizable Arrays in Optimal Time and Space



**Data blocks should
grow in size**

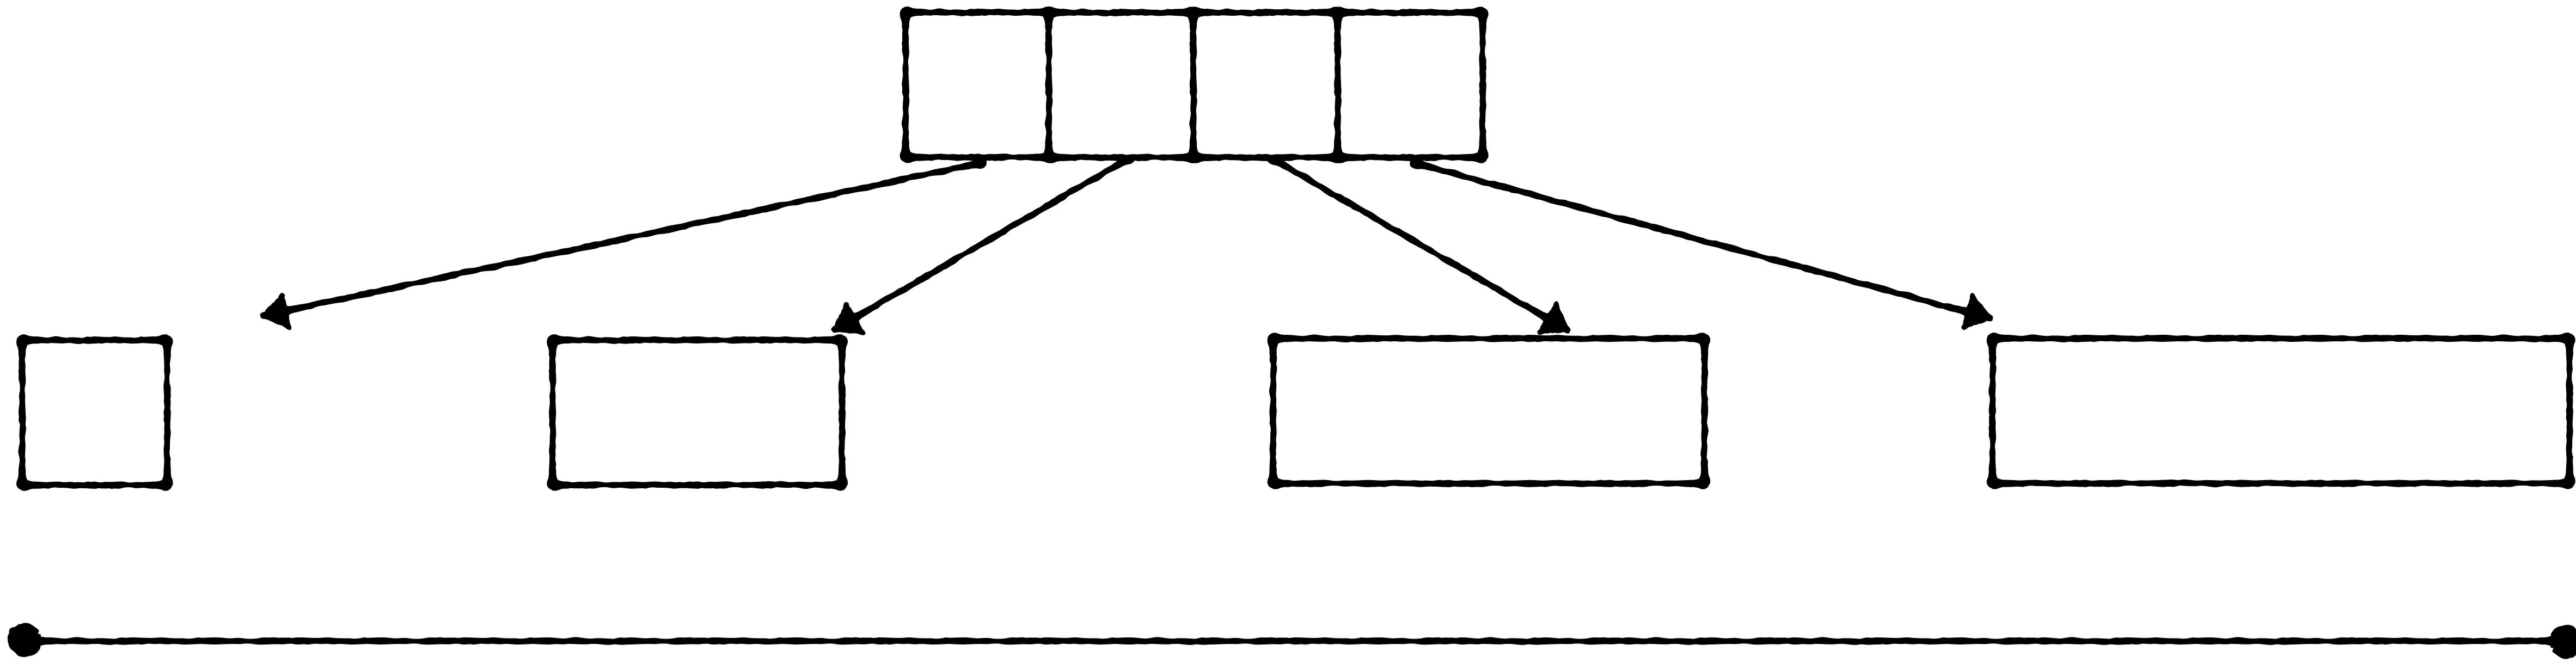
Resizable Arrays in Optimal Time and Space



Data blocks should
grow in size

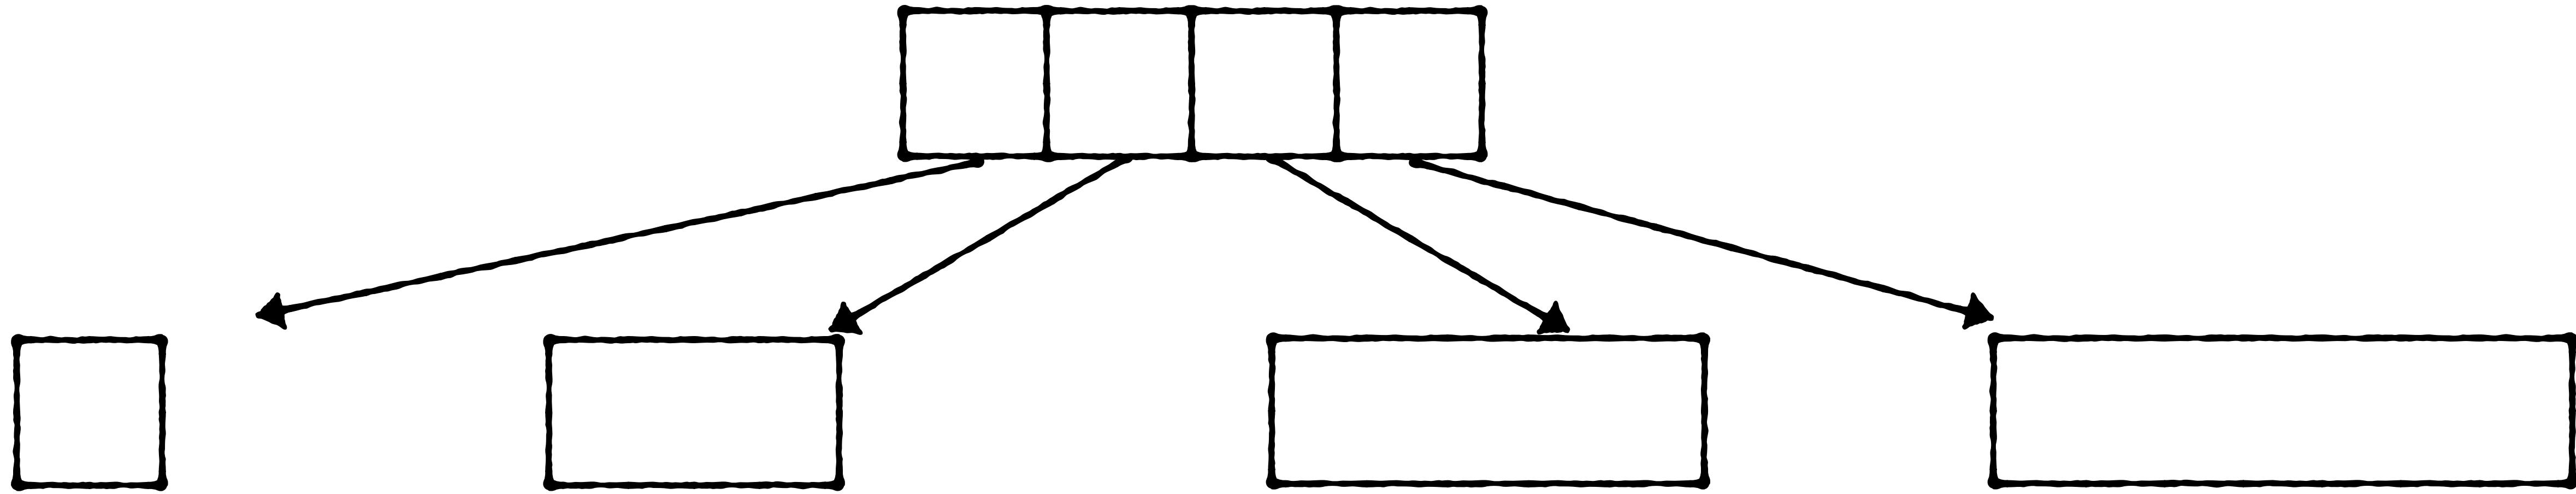


Directory grows
more slowly



$O(\sqrt{N})$ data blocks

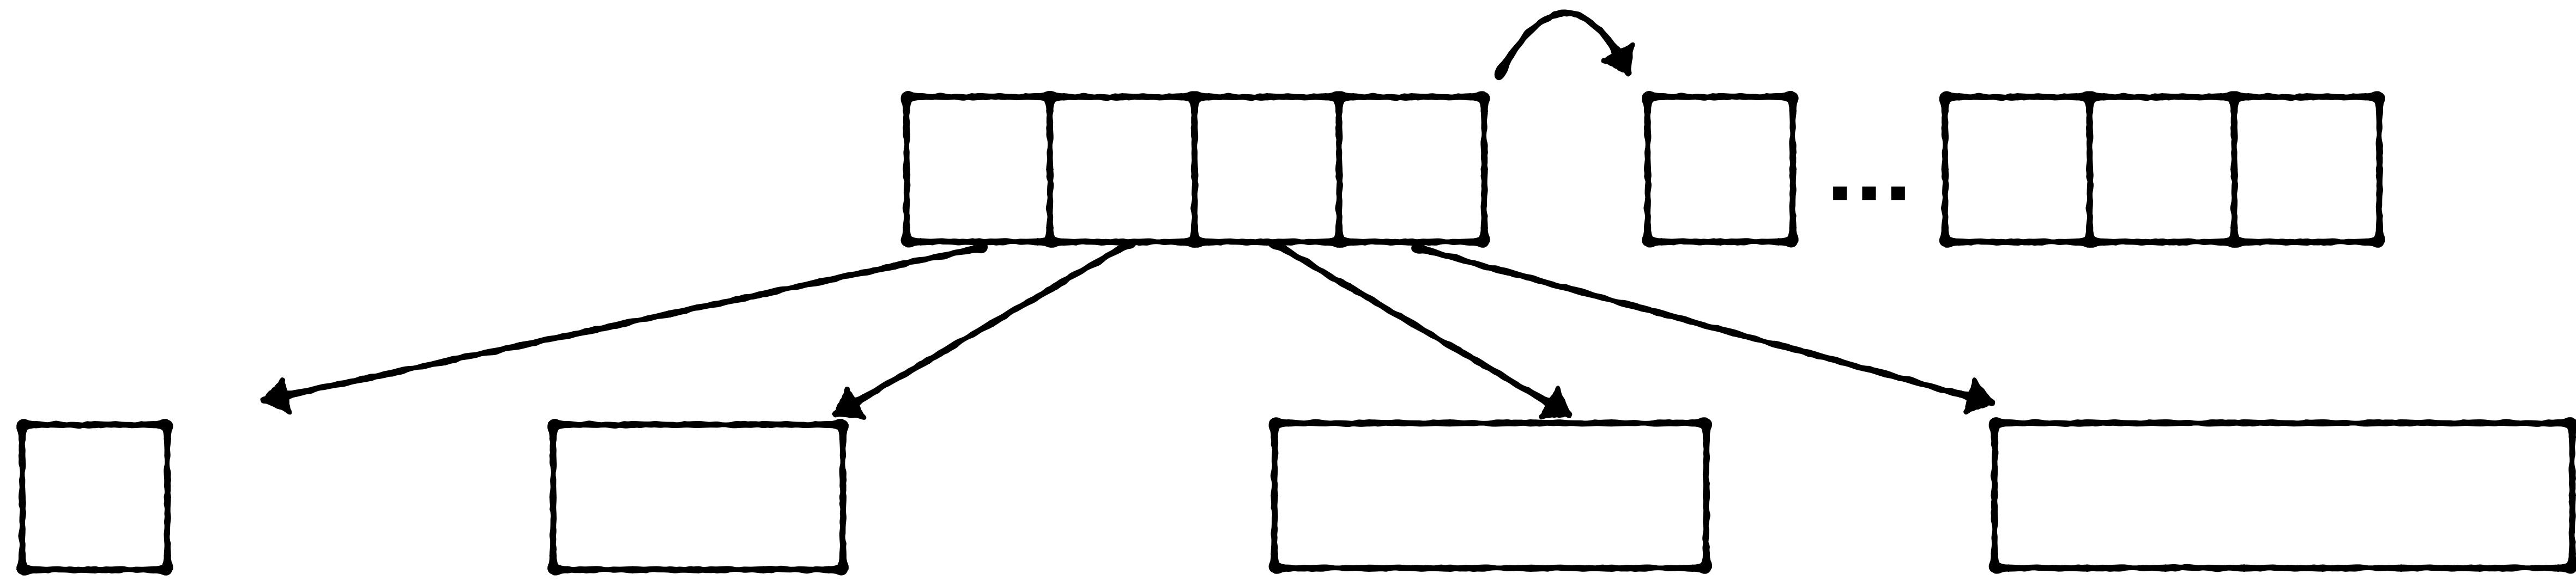
$O(\sqrt{N})$ pointers



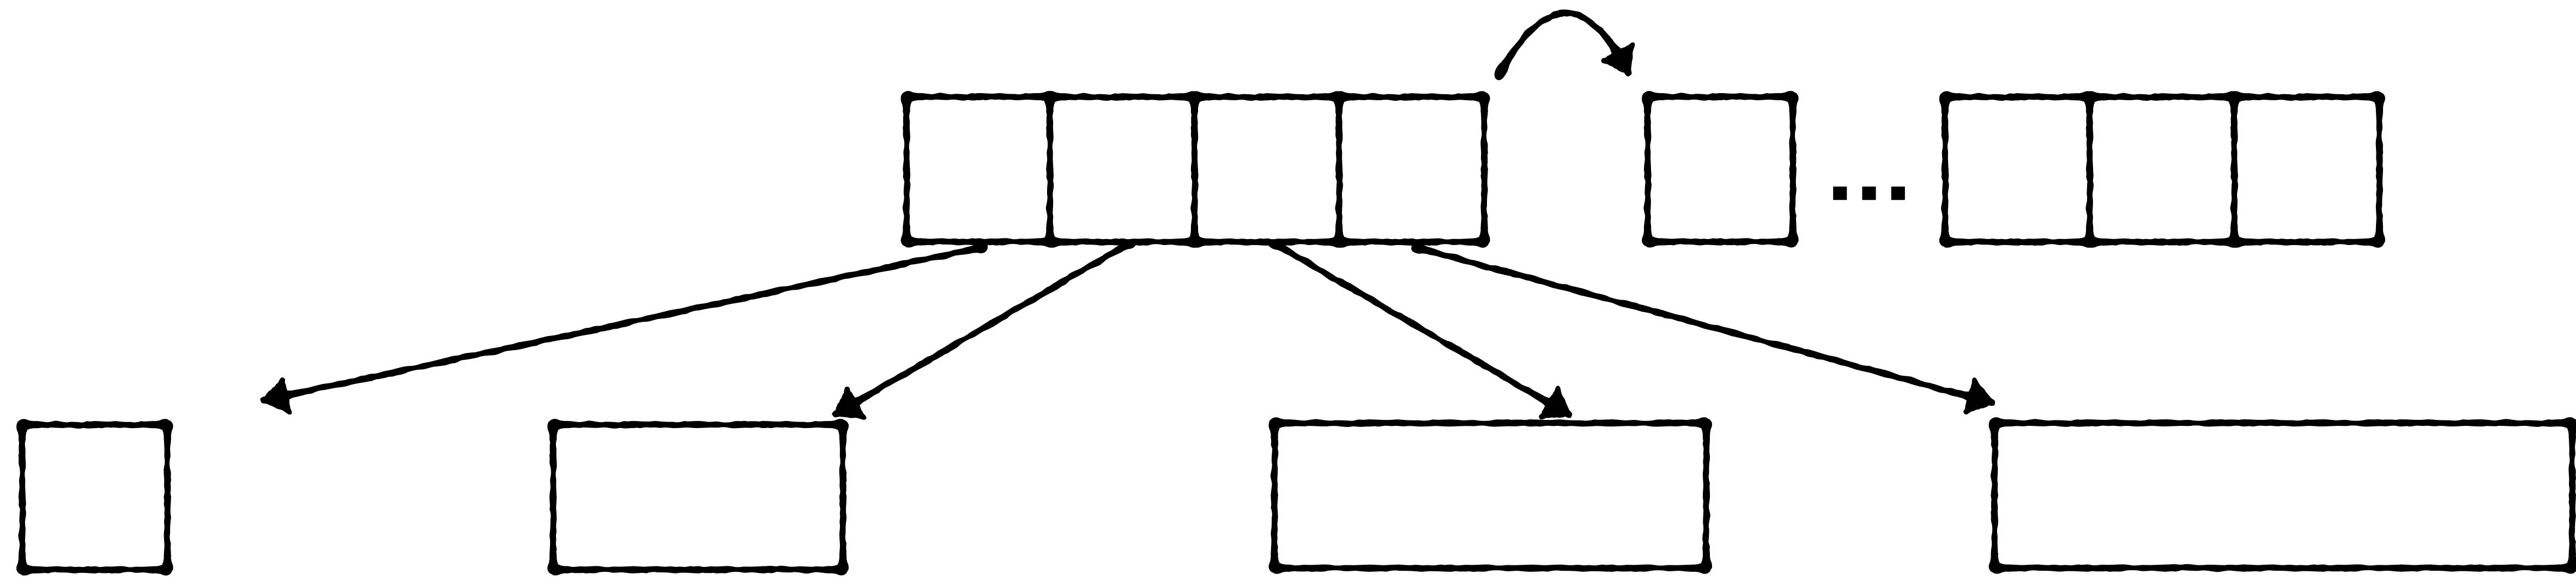
$O(\sqrt{N})$ data blocks

$O(\sqrt{N})$ pointers

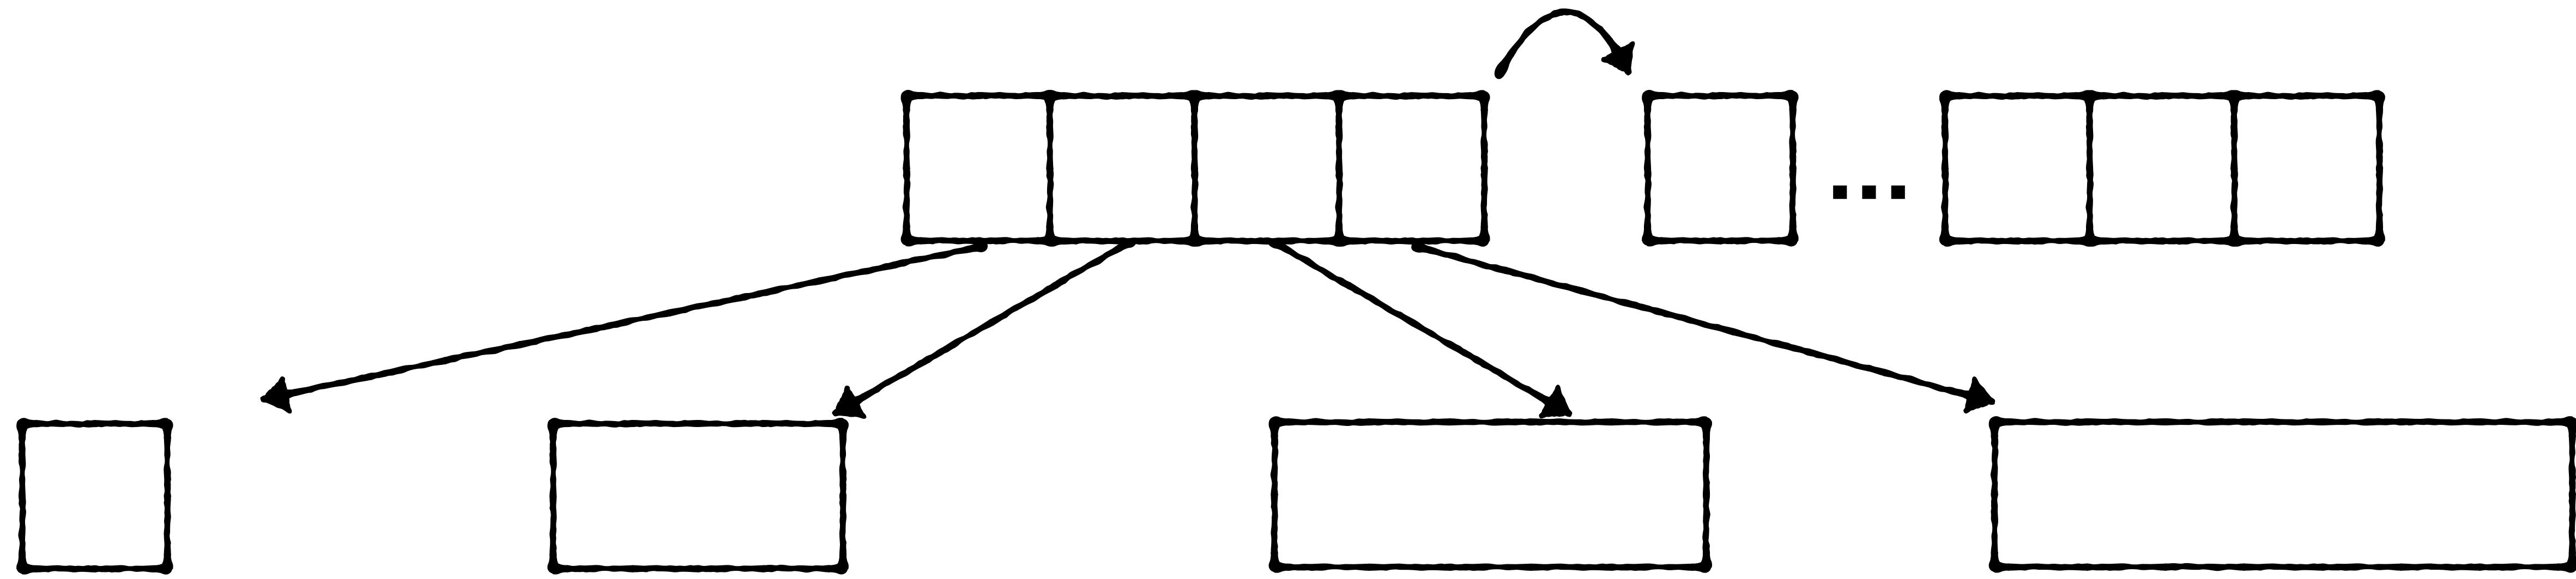
2x when full



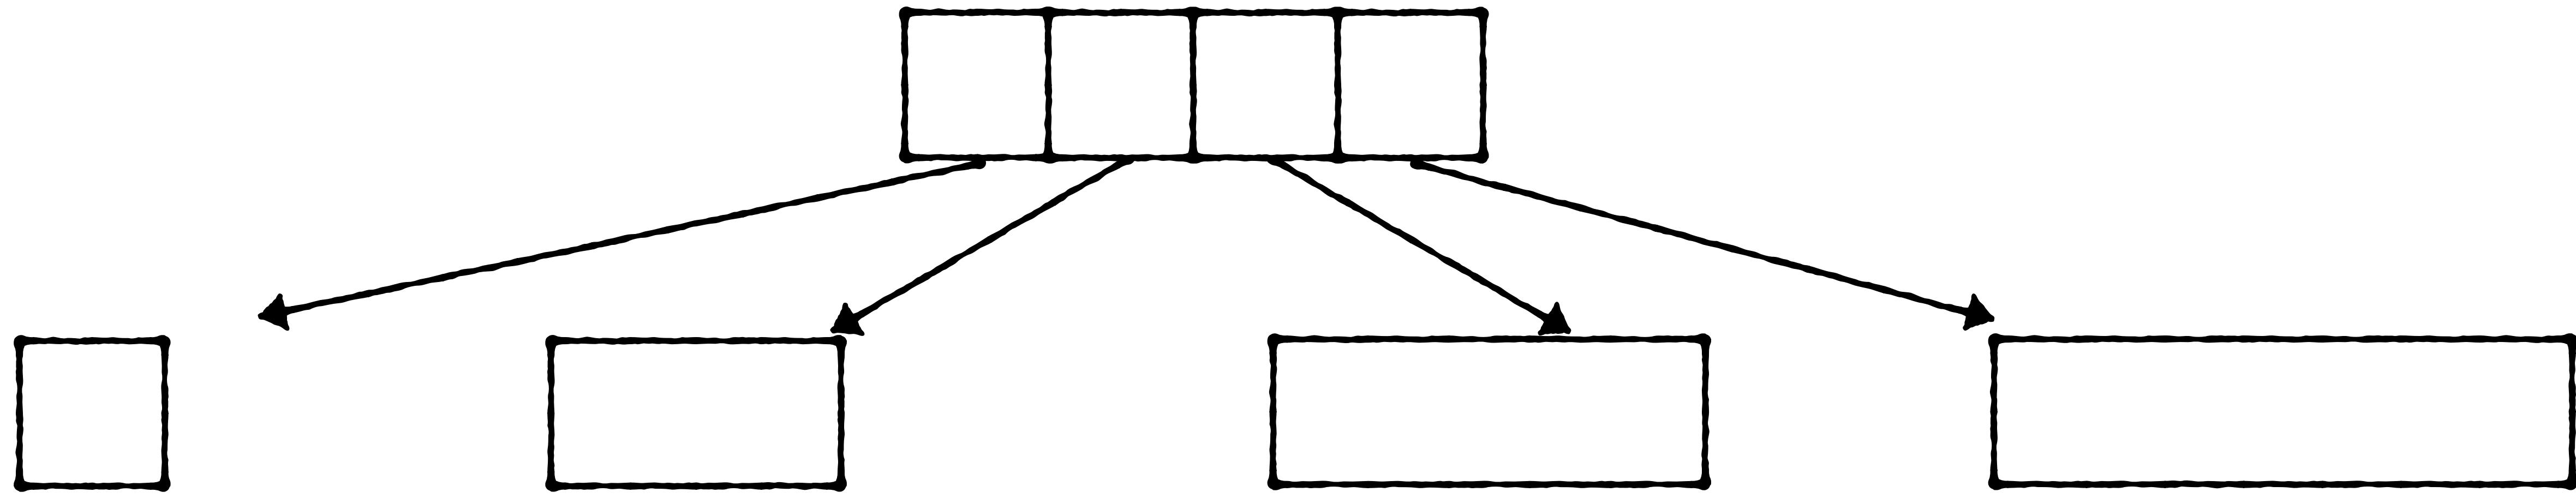
$O(2\sqrt{N})$ pointers



$O(\sqrt{N})$ pointers

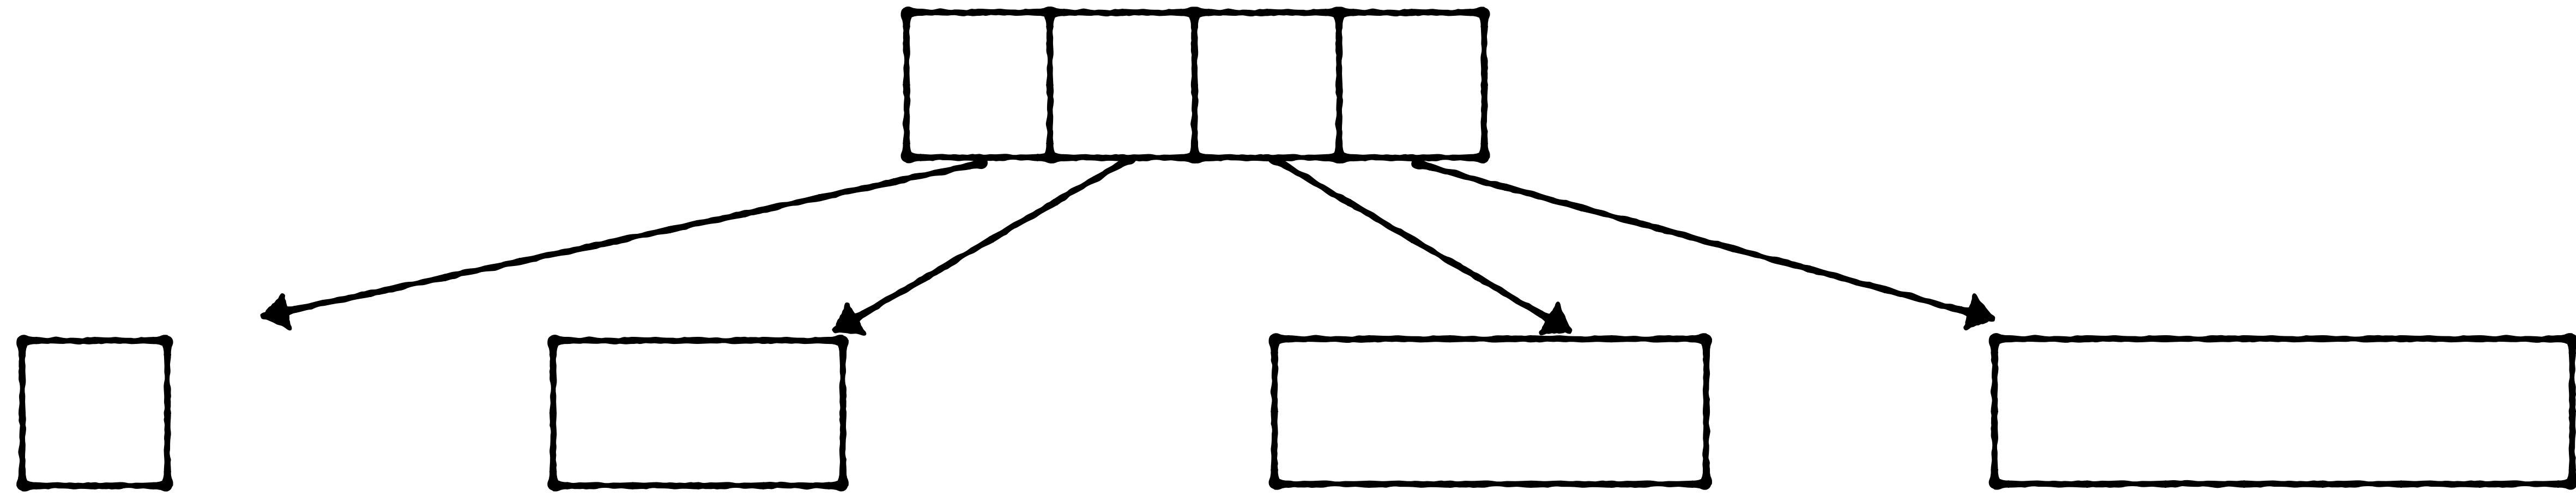


$O(\sqrt{N})$ pointers



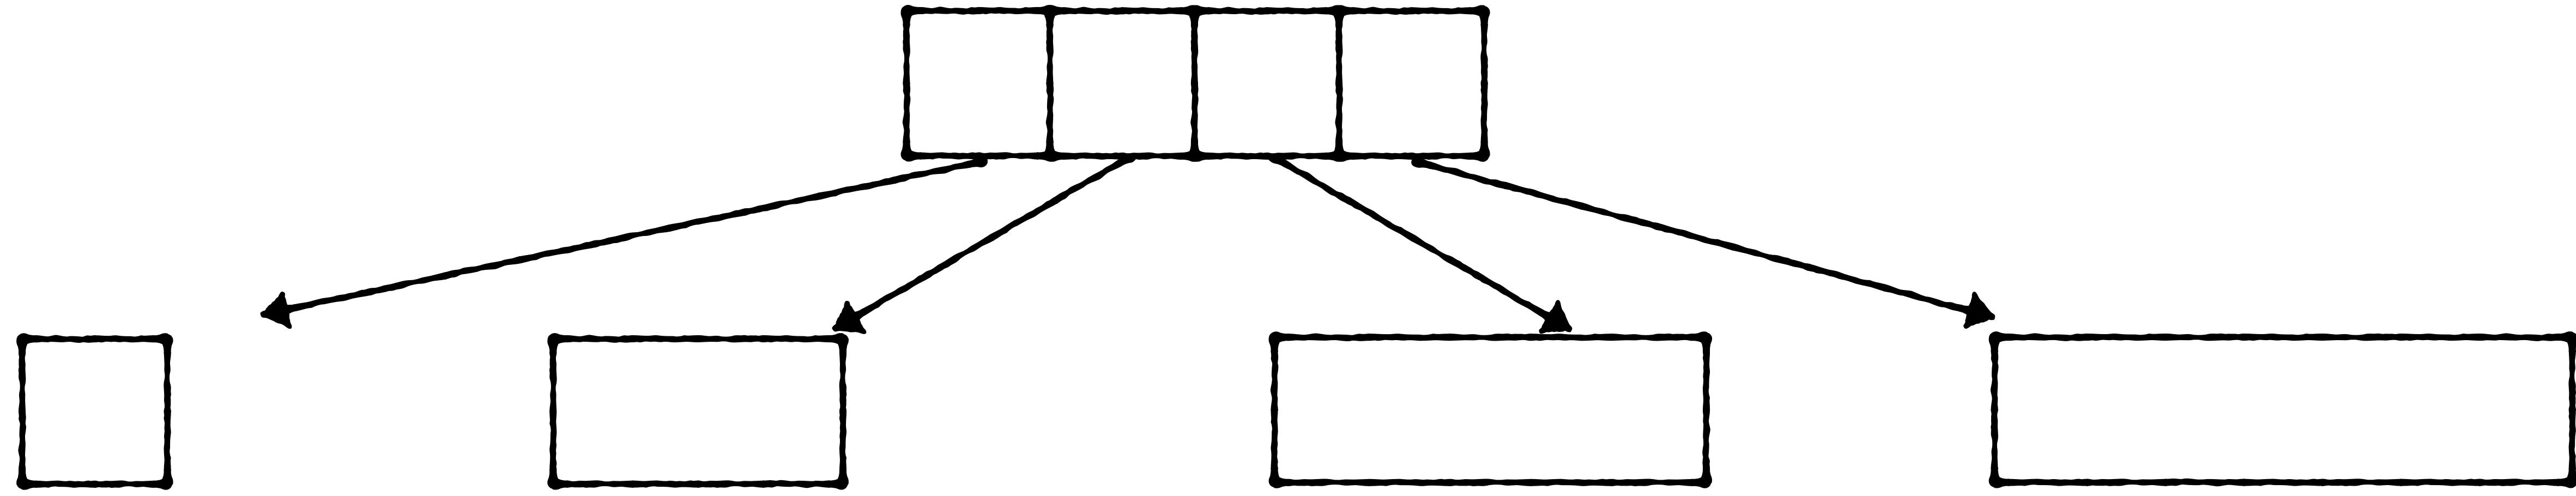
$O(\sqrt{N})$ slots

$O(\sqrt{N})$ pointers

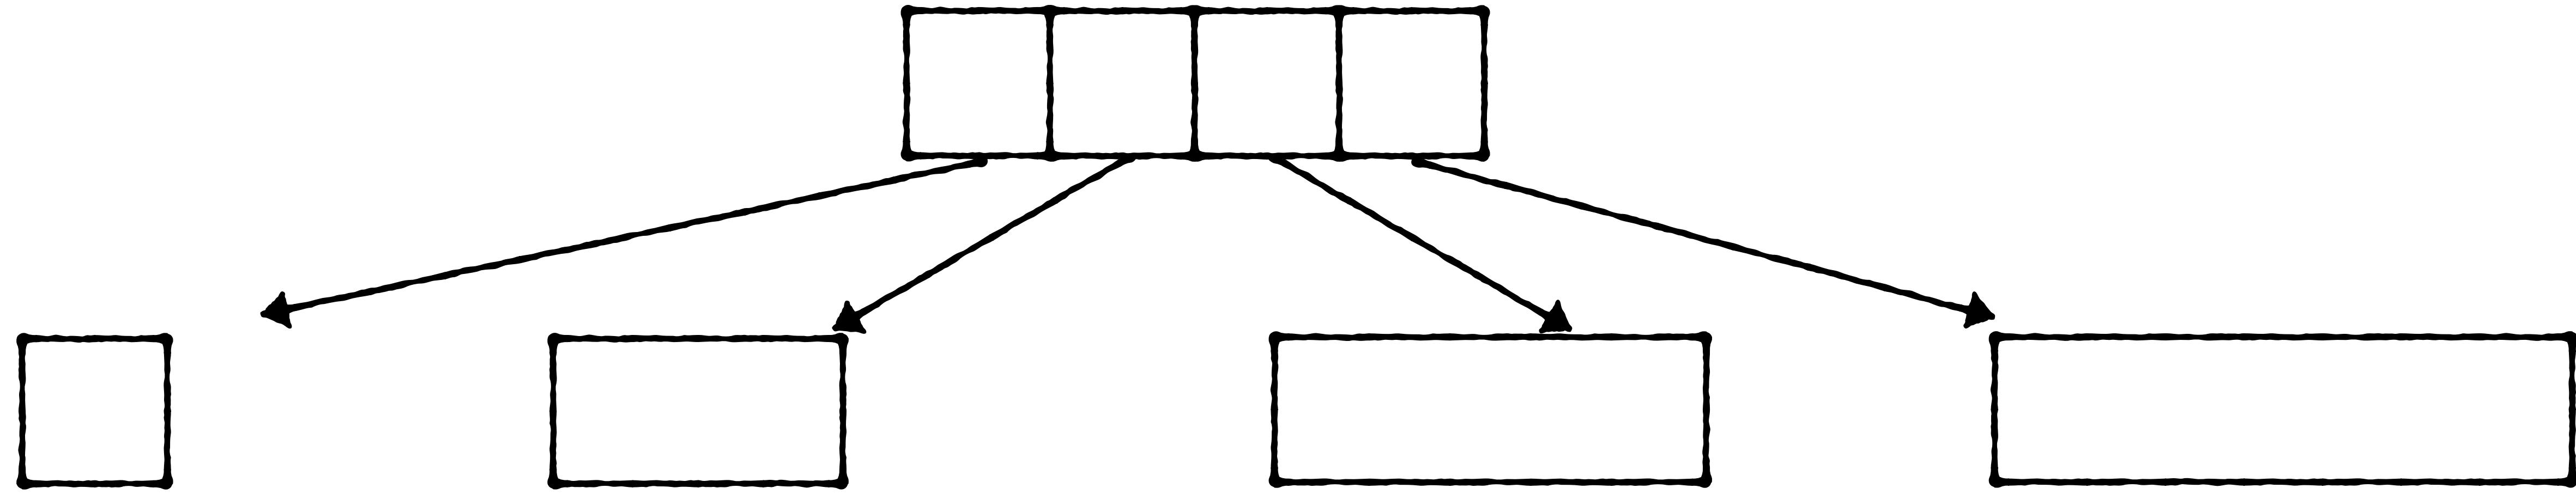


**Waste at most
 $O(\sqrt{N})$ slots**

Max space amp = $O(\sqrt{N}) + O(\sqrt{N}) = O(\sqrt{N})$

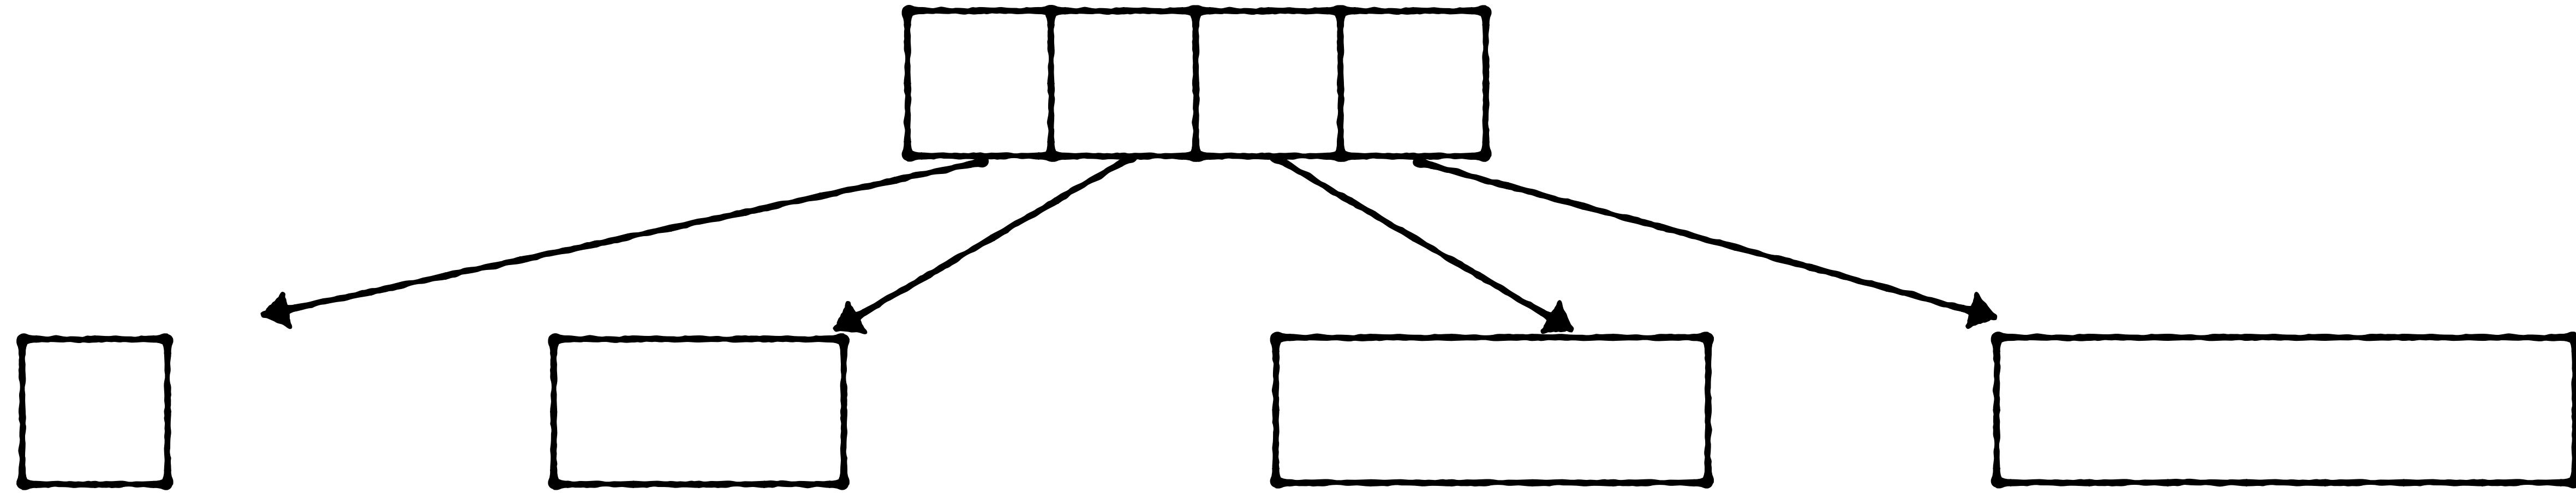


Max space amp = $O(\sqrt{N})$



Challenges: How to grow blocks to meet these properties?

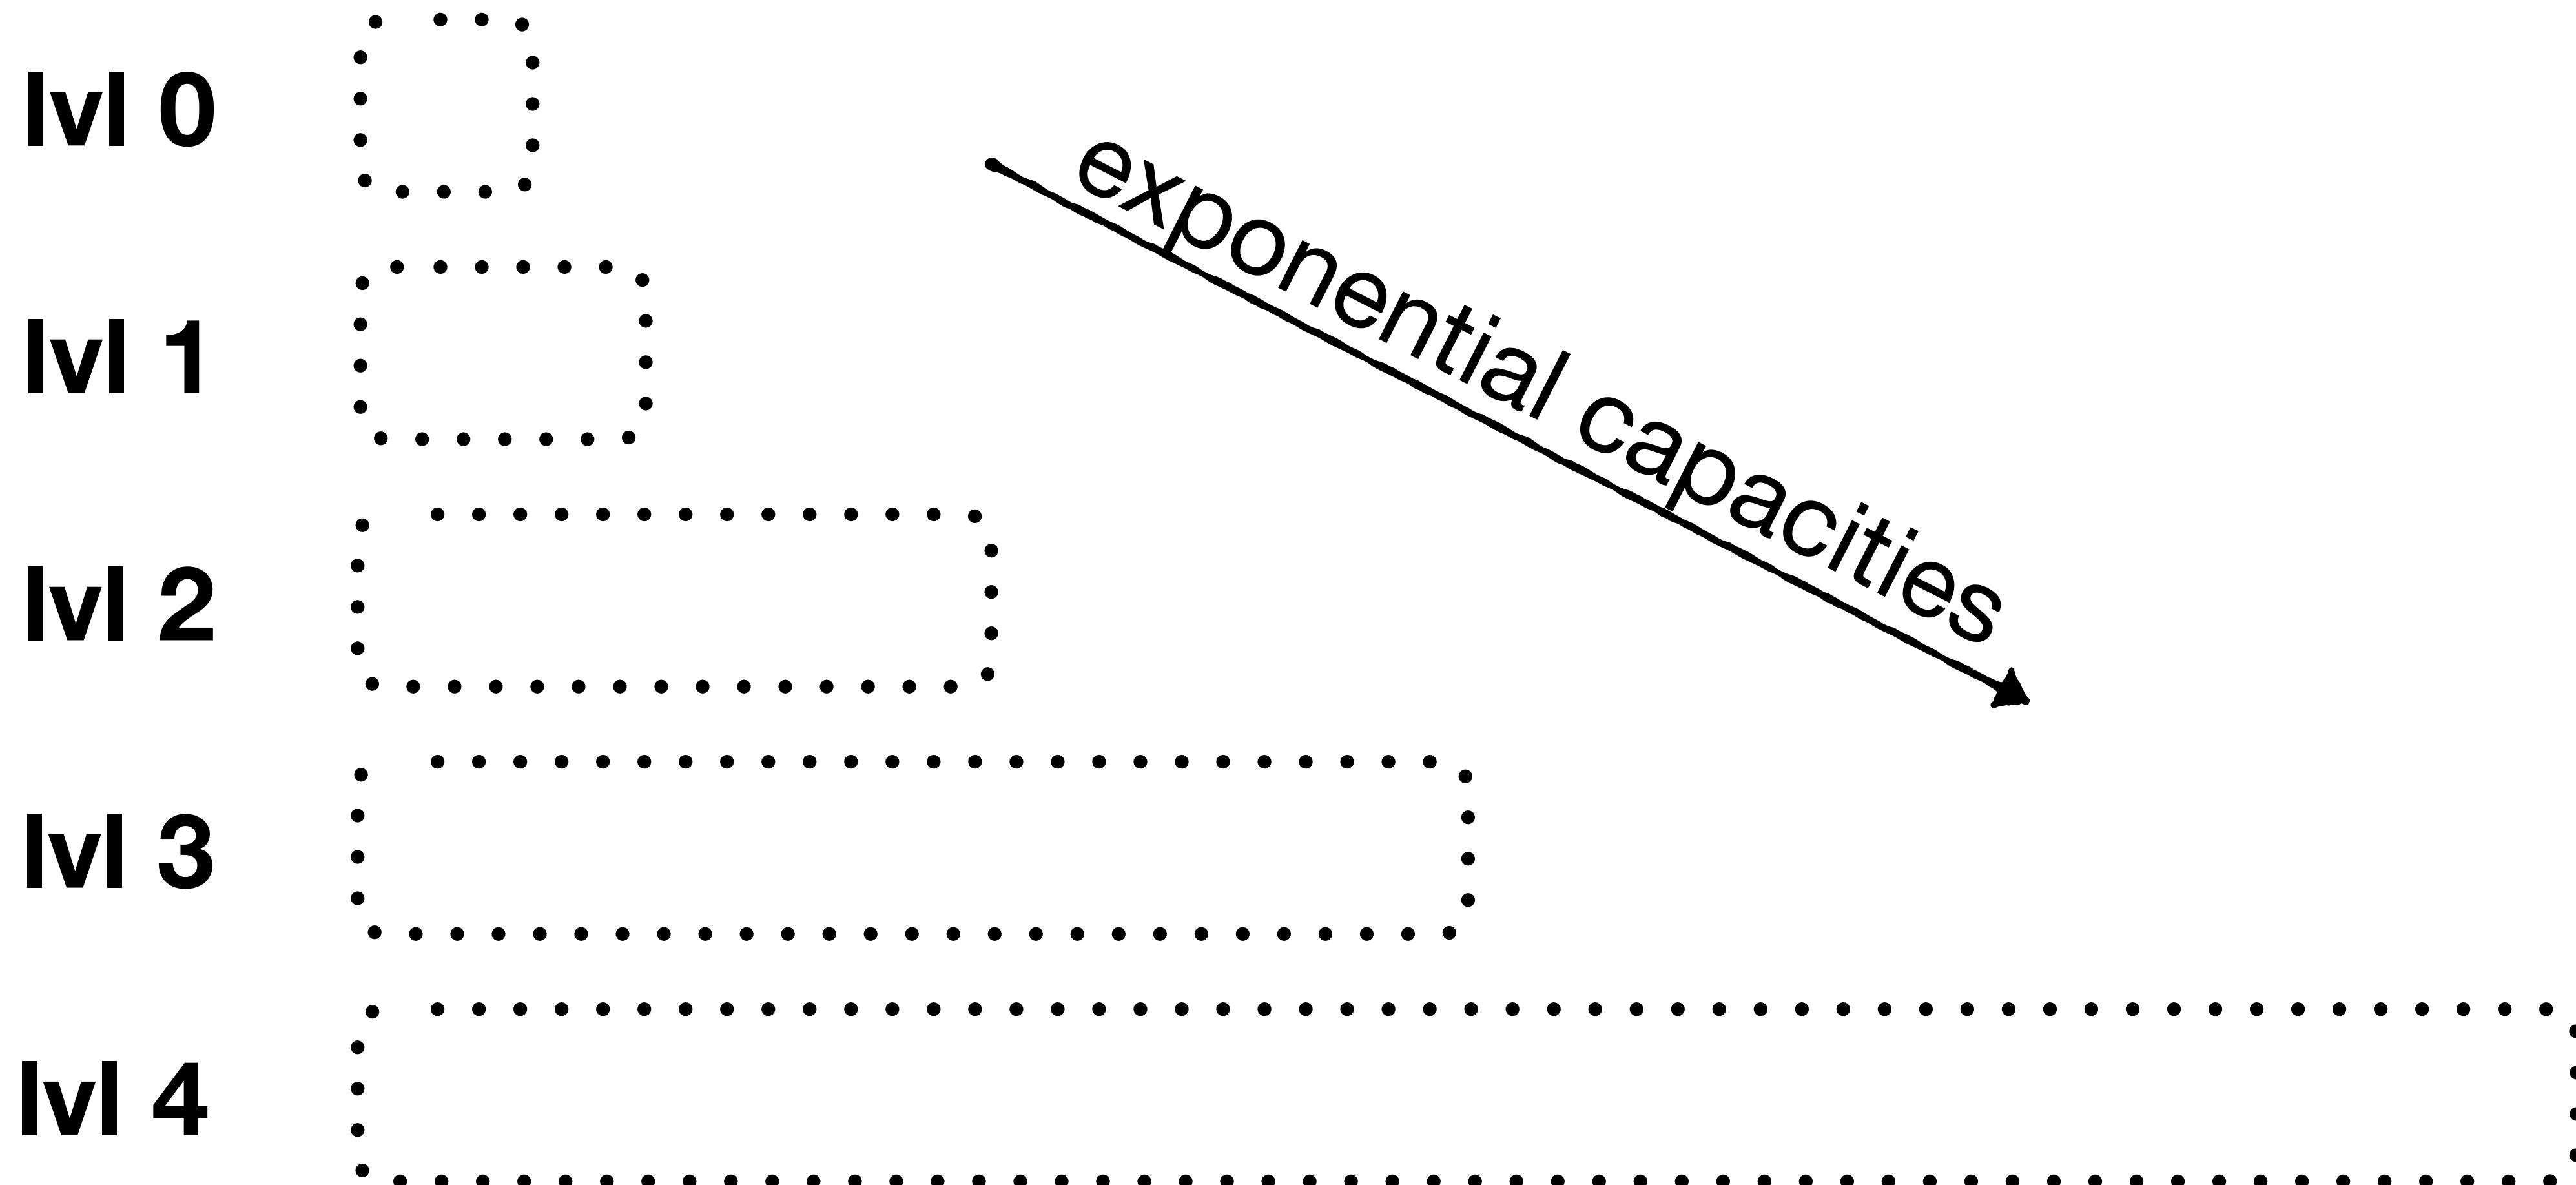
Max space amp = $O(\sqrt{N})$



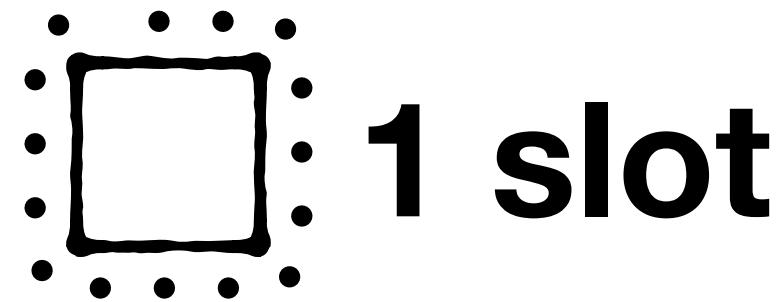
Challenges: How to grow blocks to meet these properties?

Inferring which block contains which array offset?

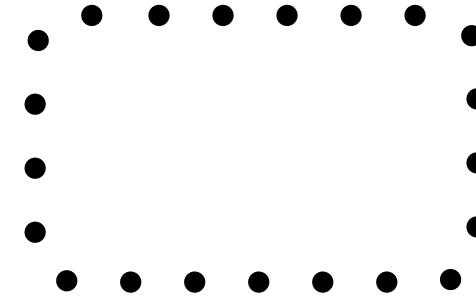
Multiple levels



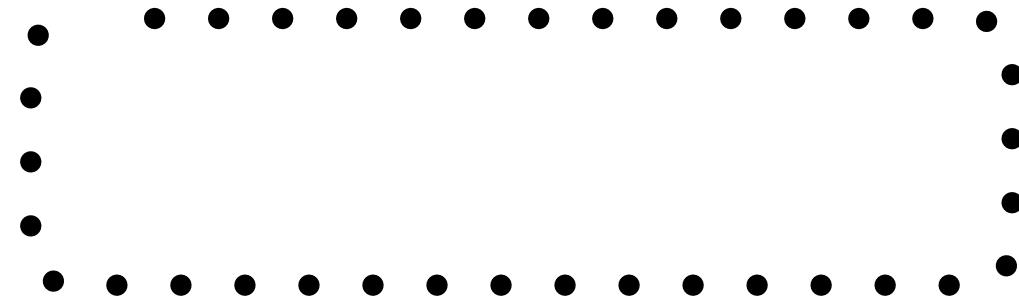
lvl 0



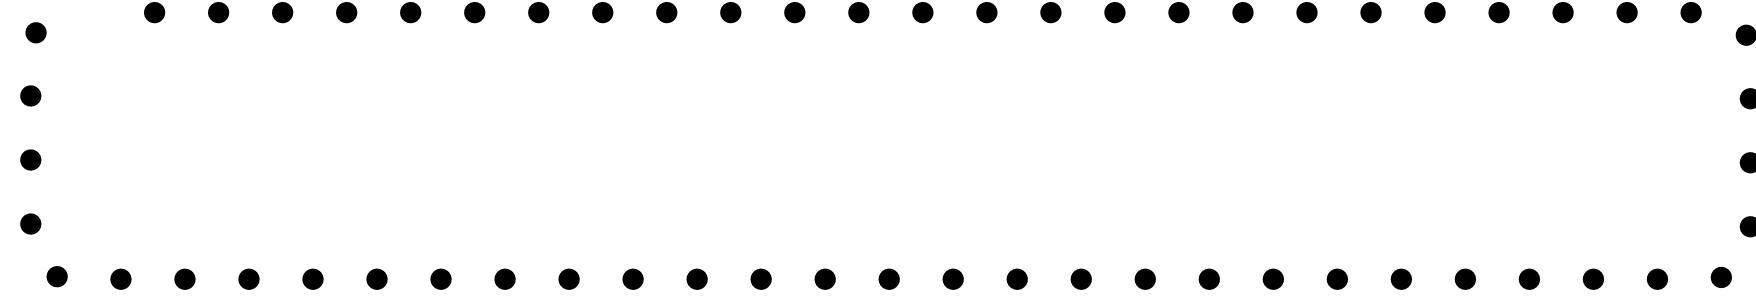
lvl 1



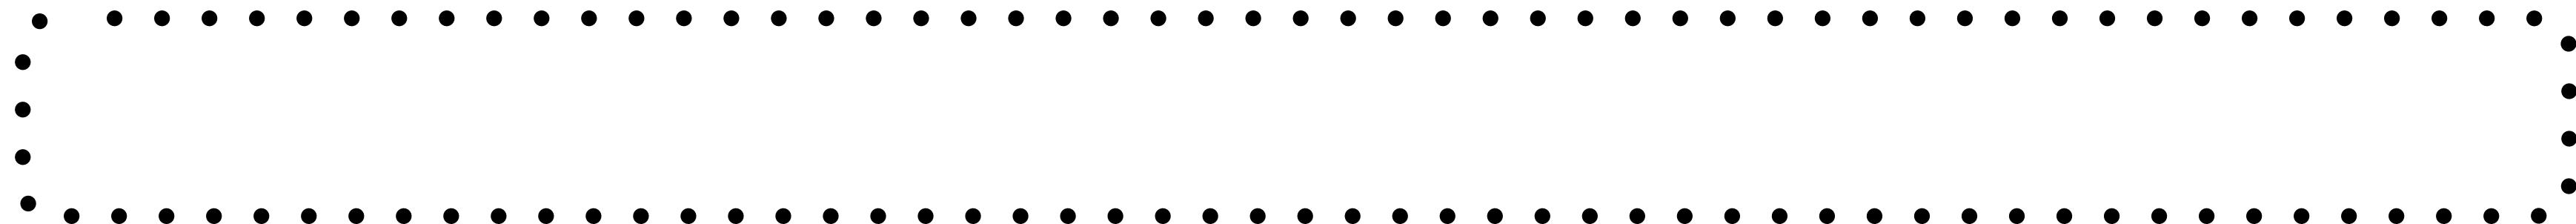
lvl 2



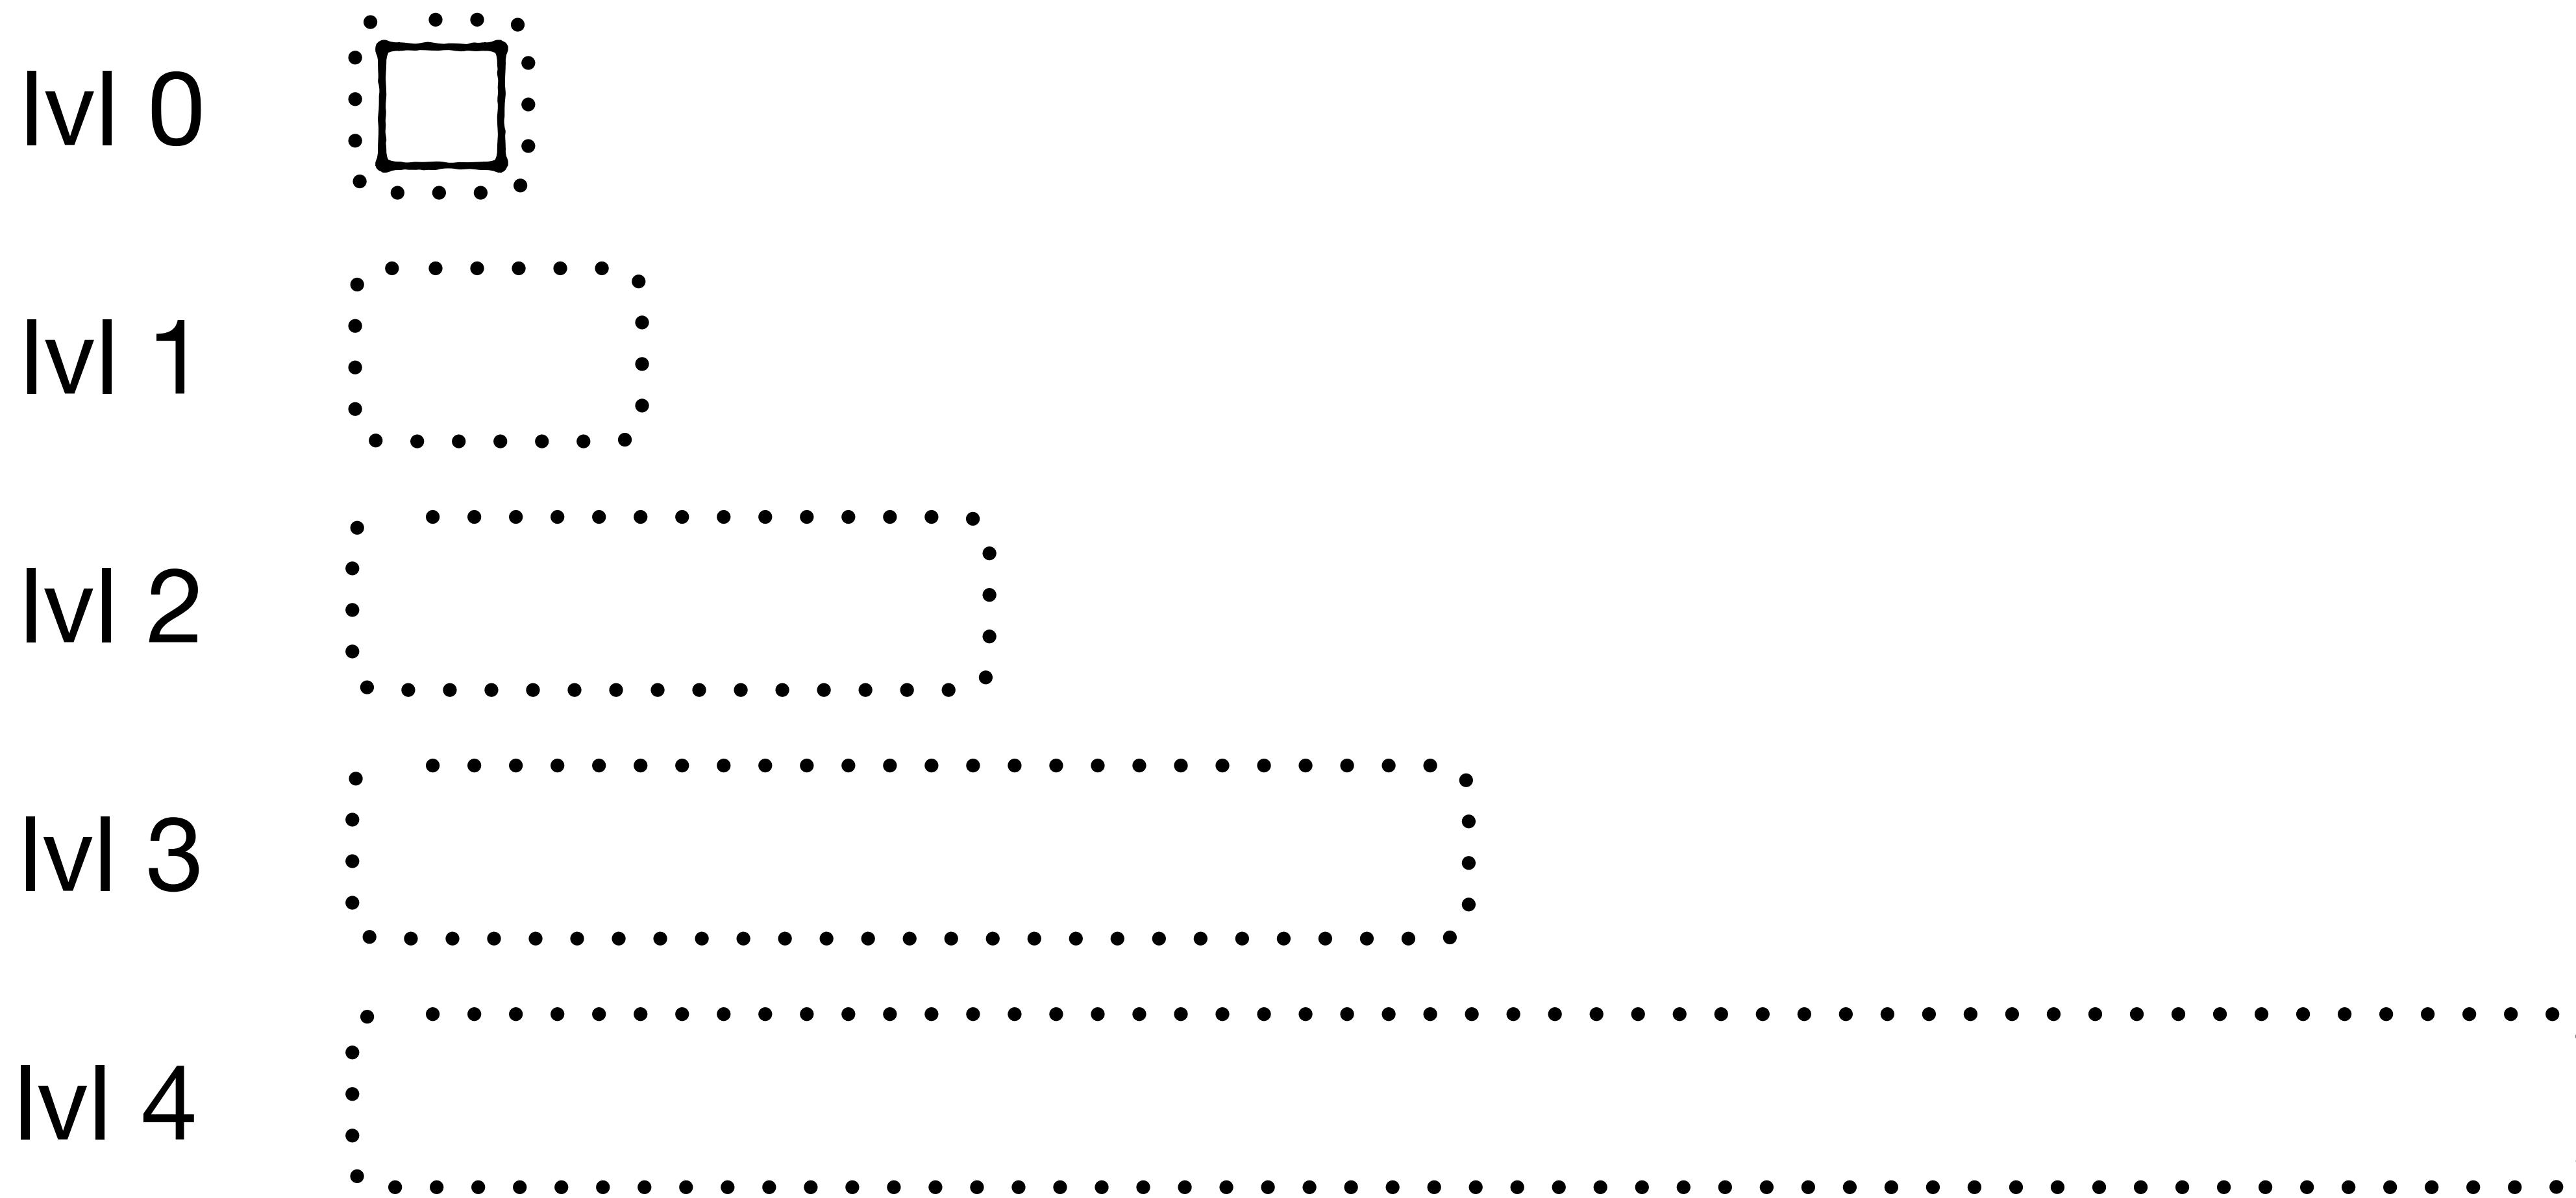
lvl 3



lvl 4

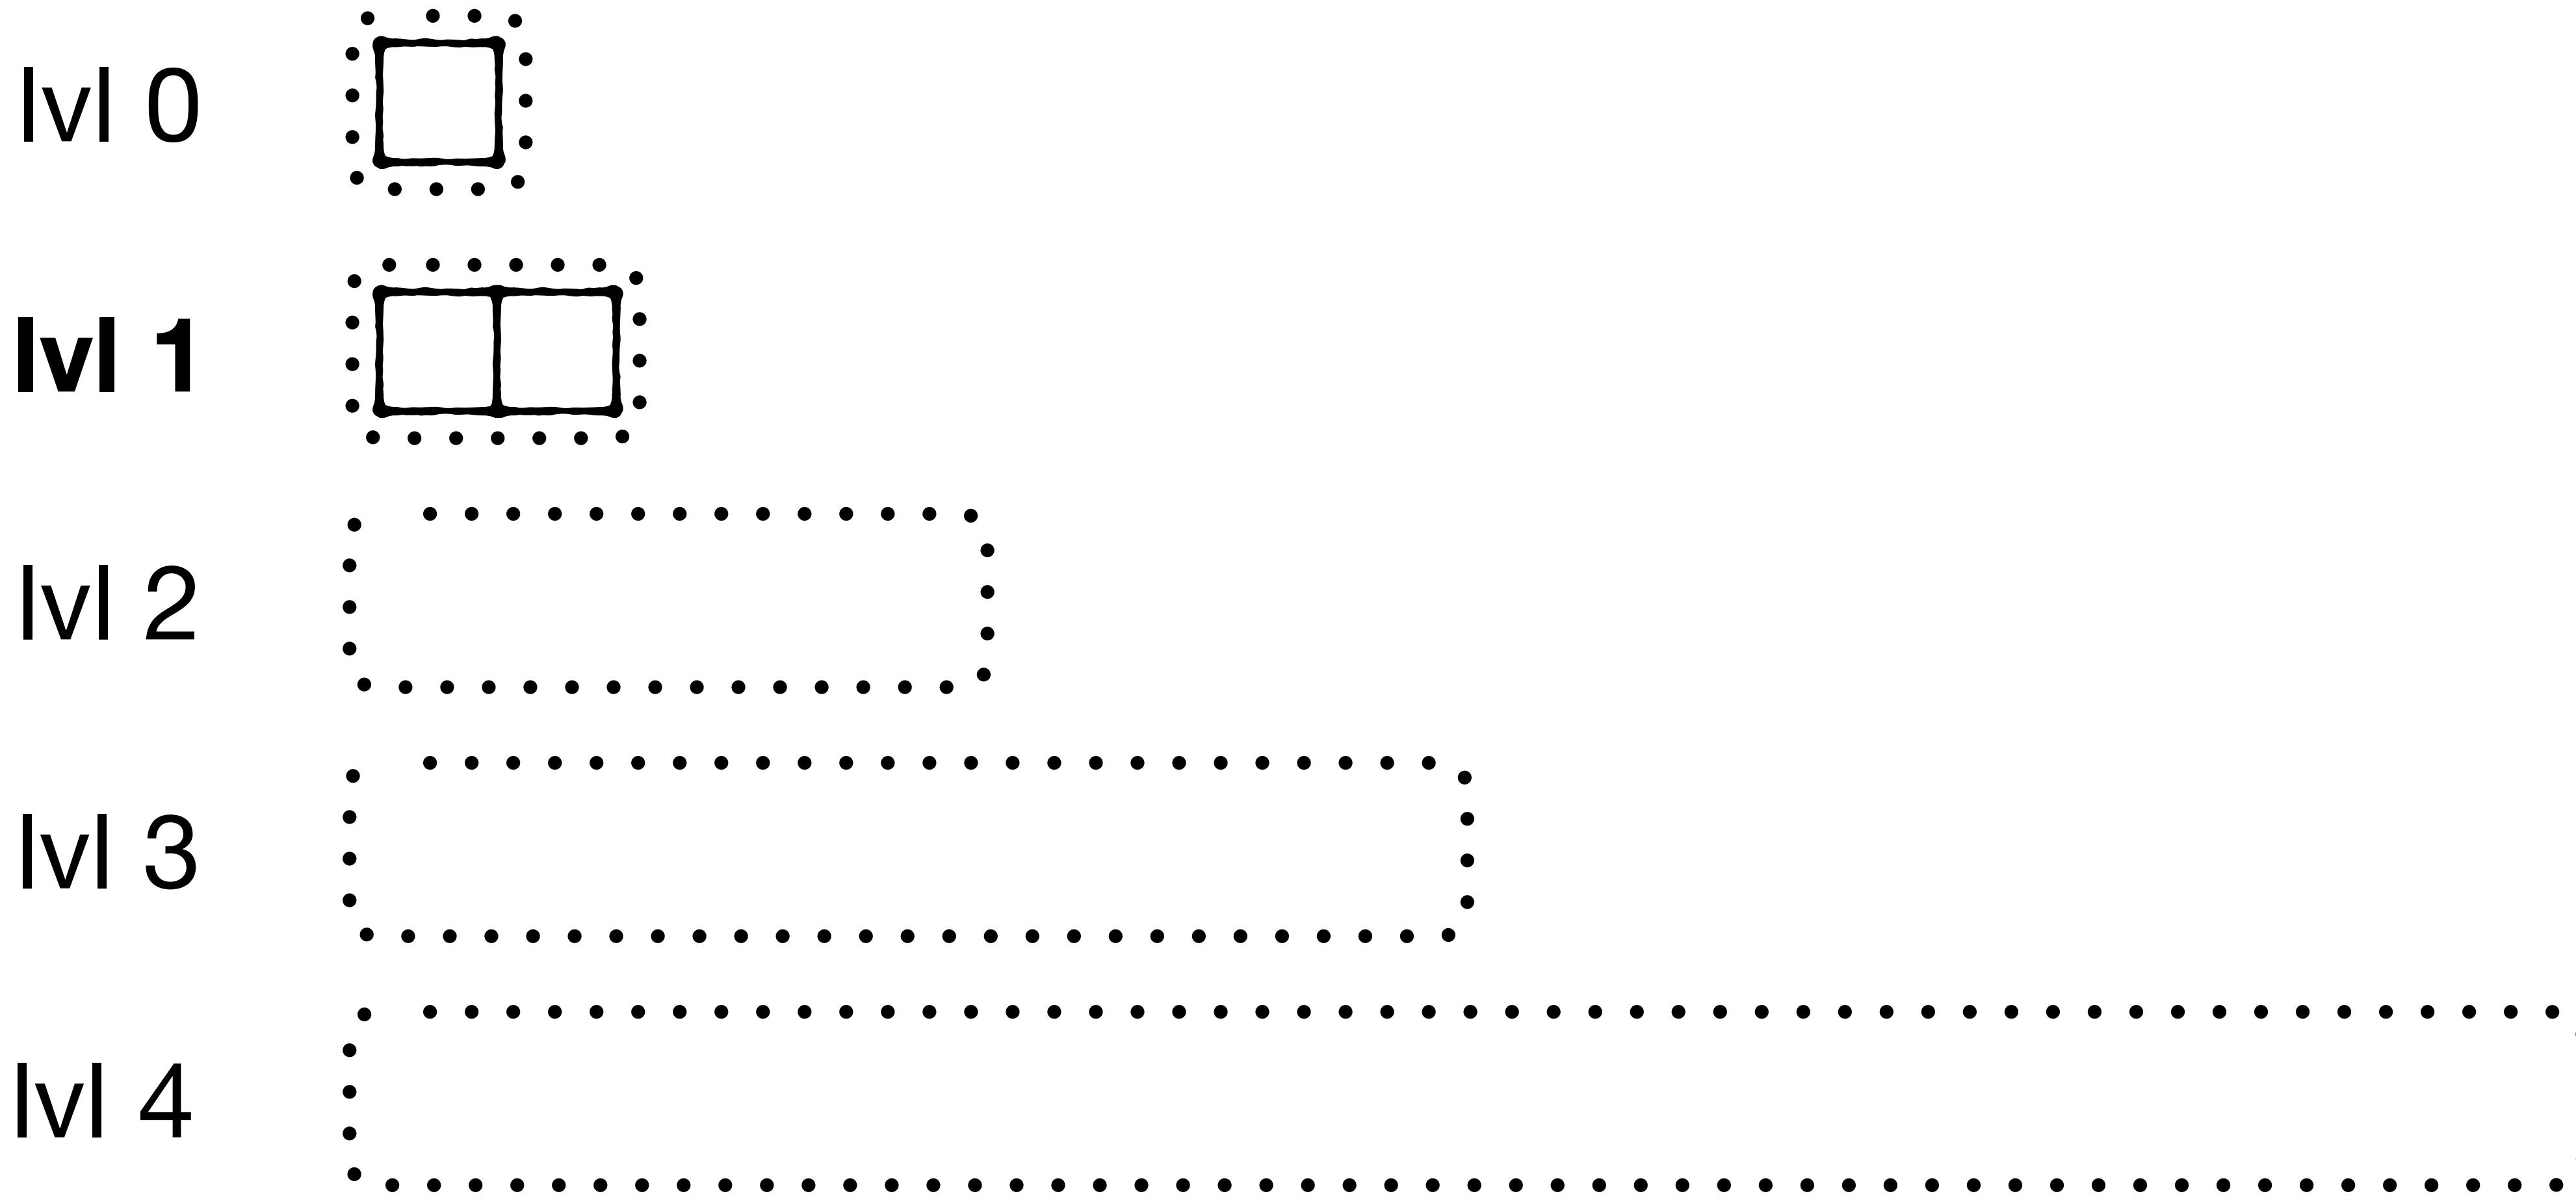


In every pair of subsequent levels k and $k+1$



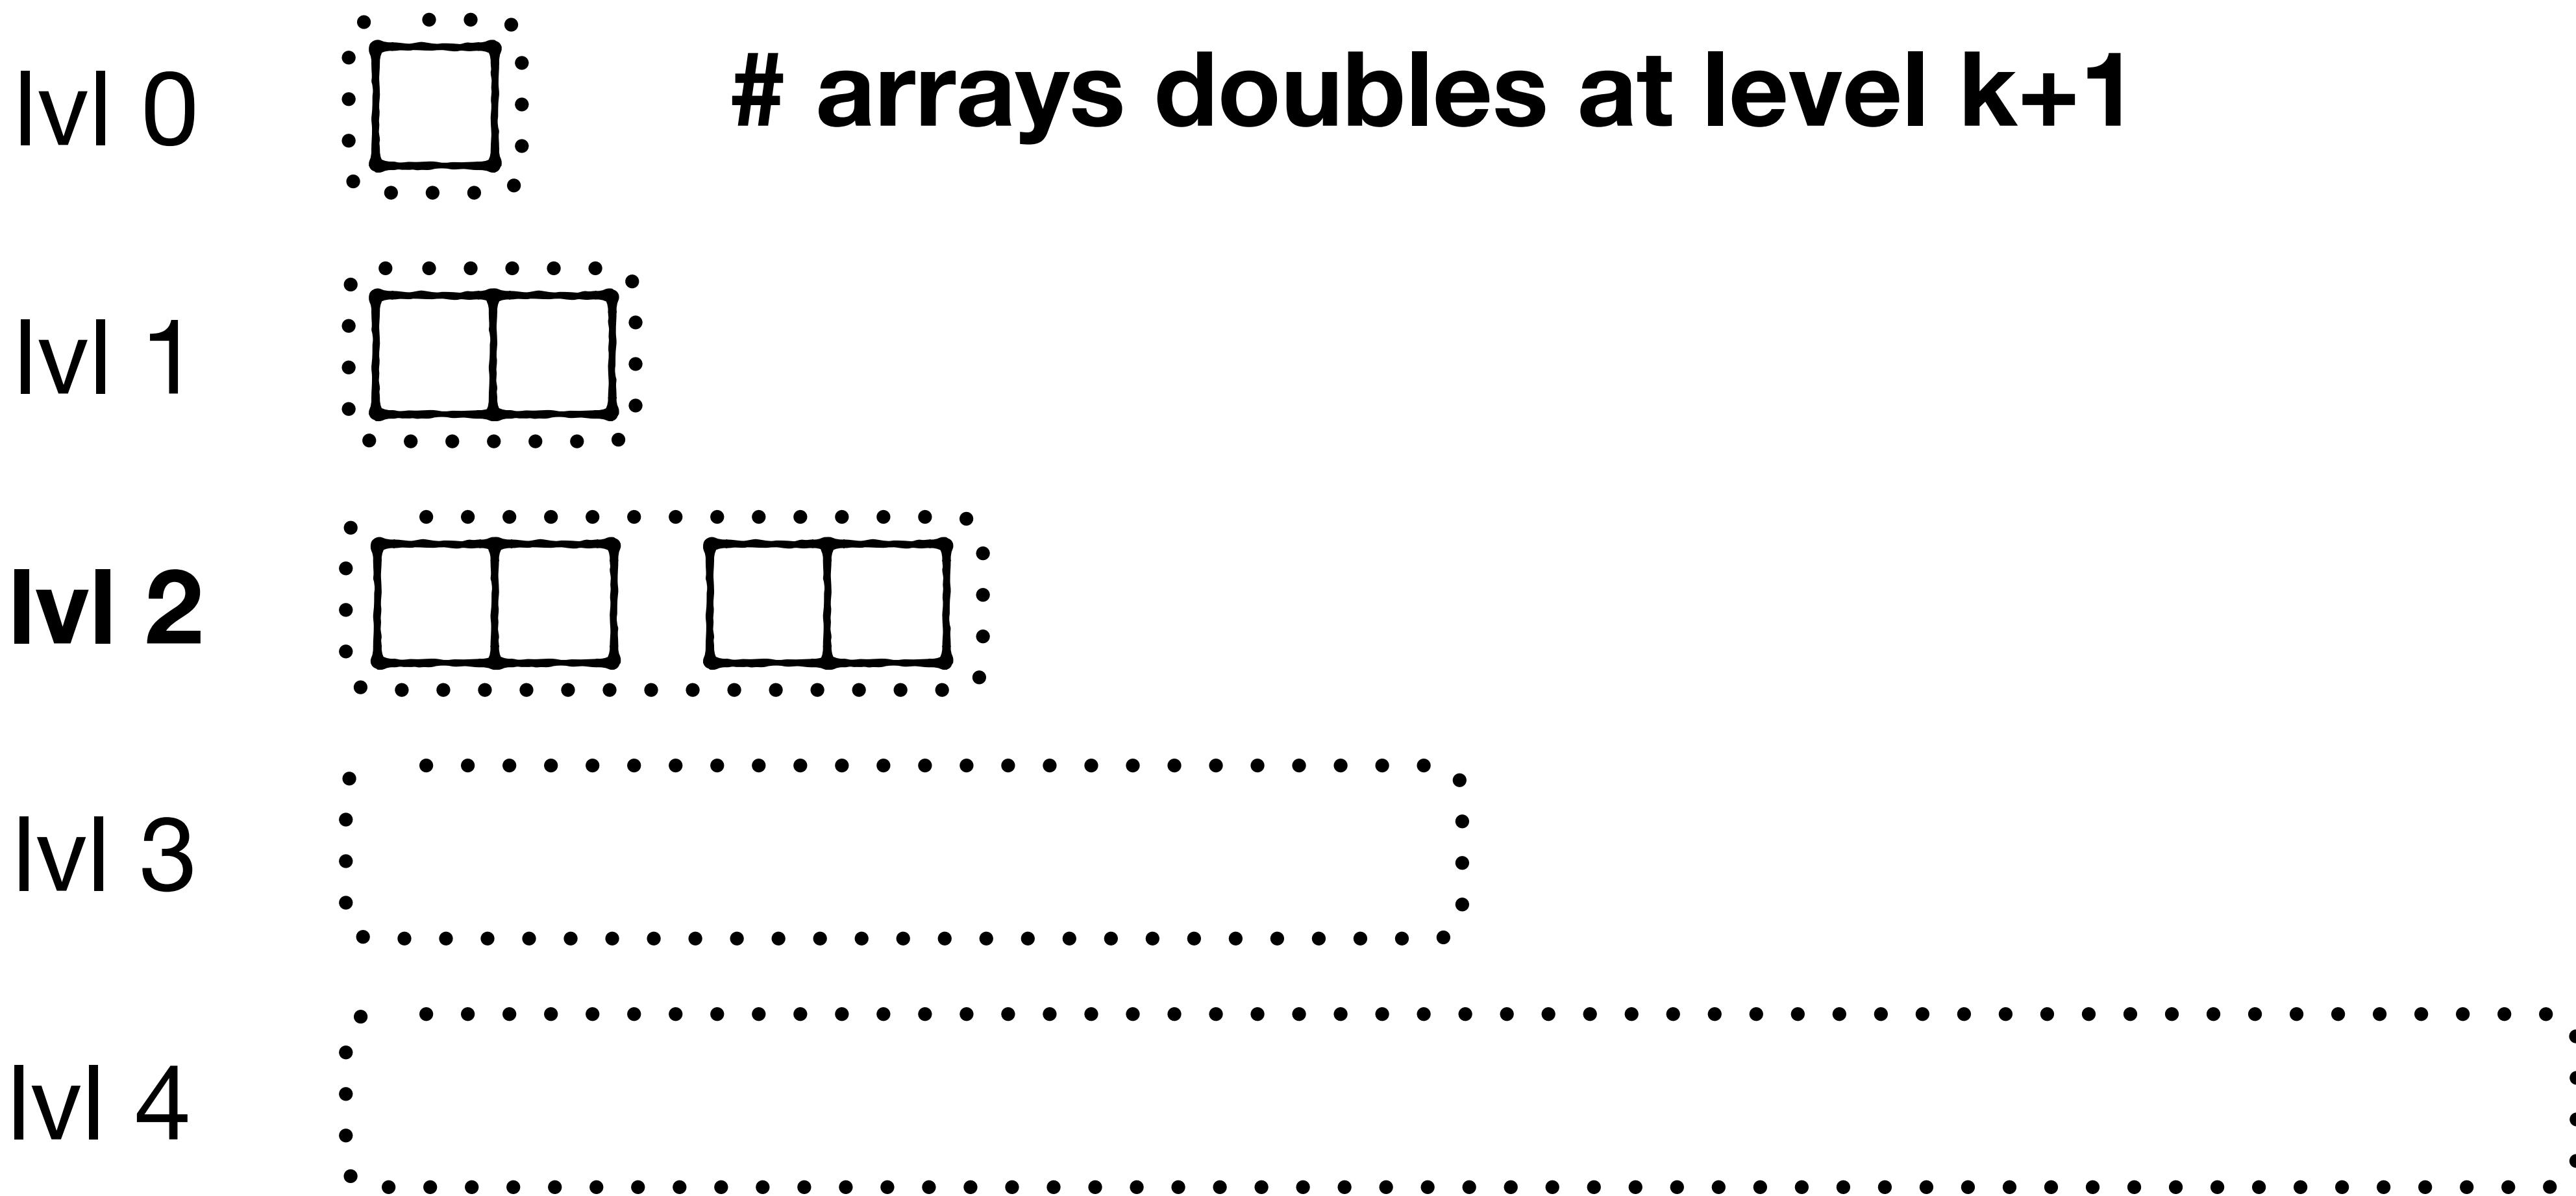
In every pair of subsequent levels k and $k+1$

Size of arrays doubles at level k



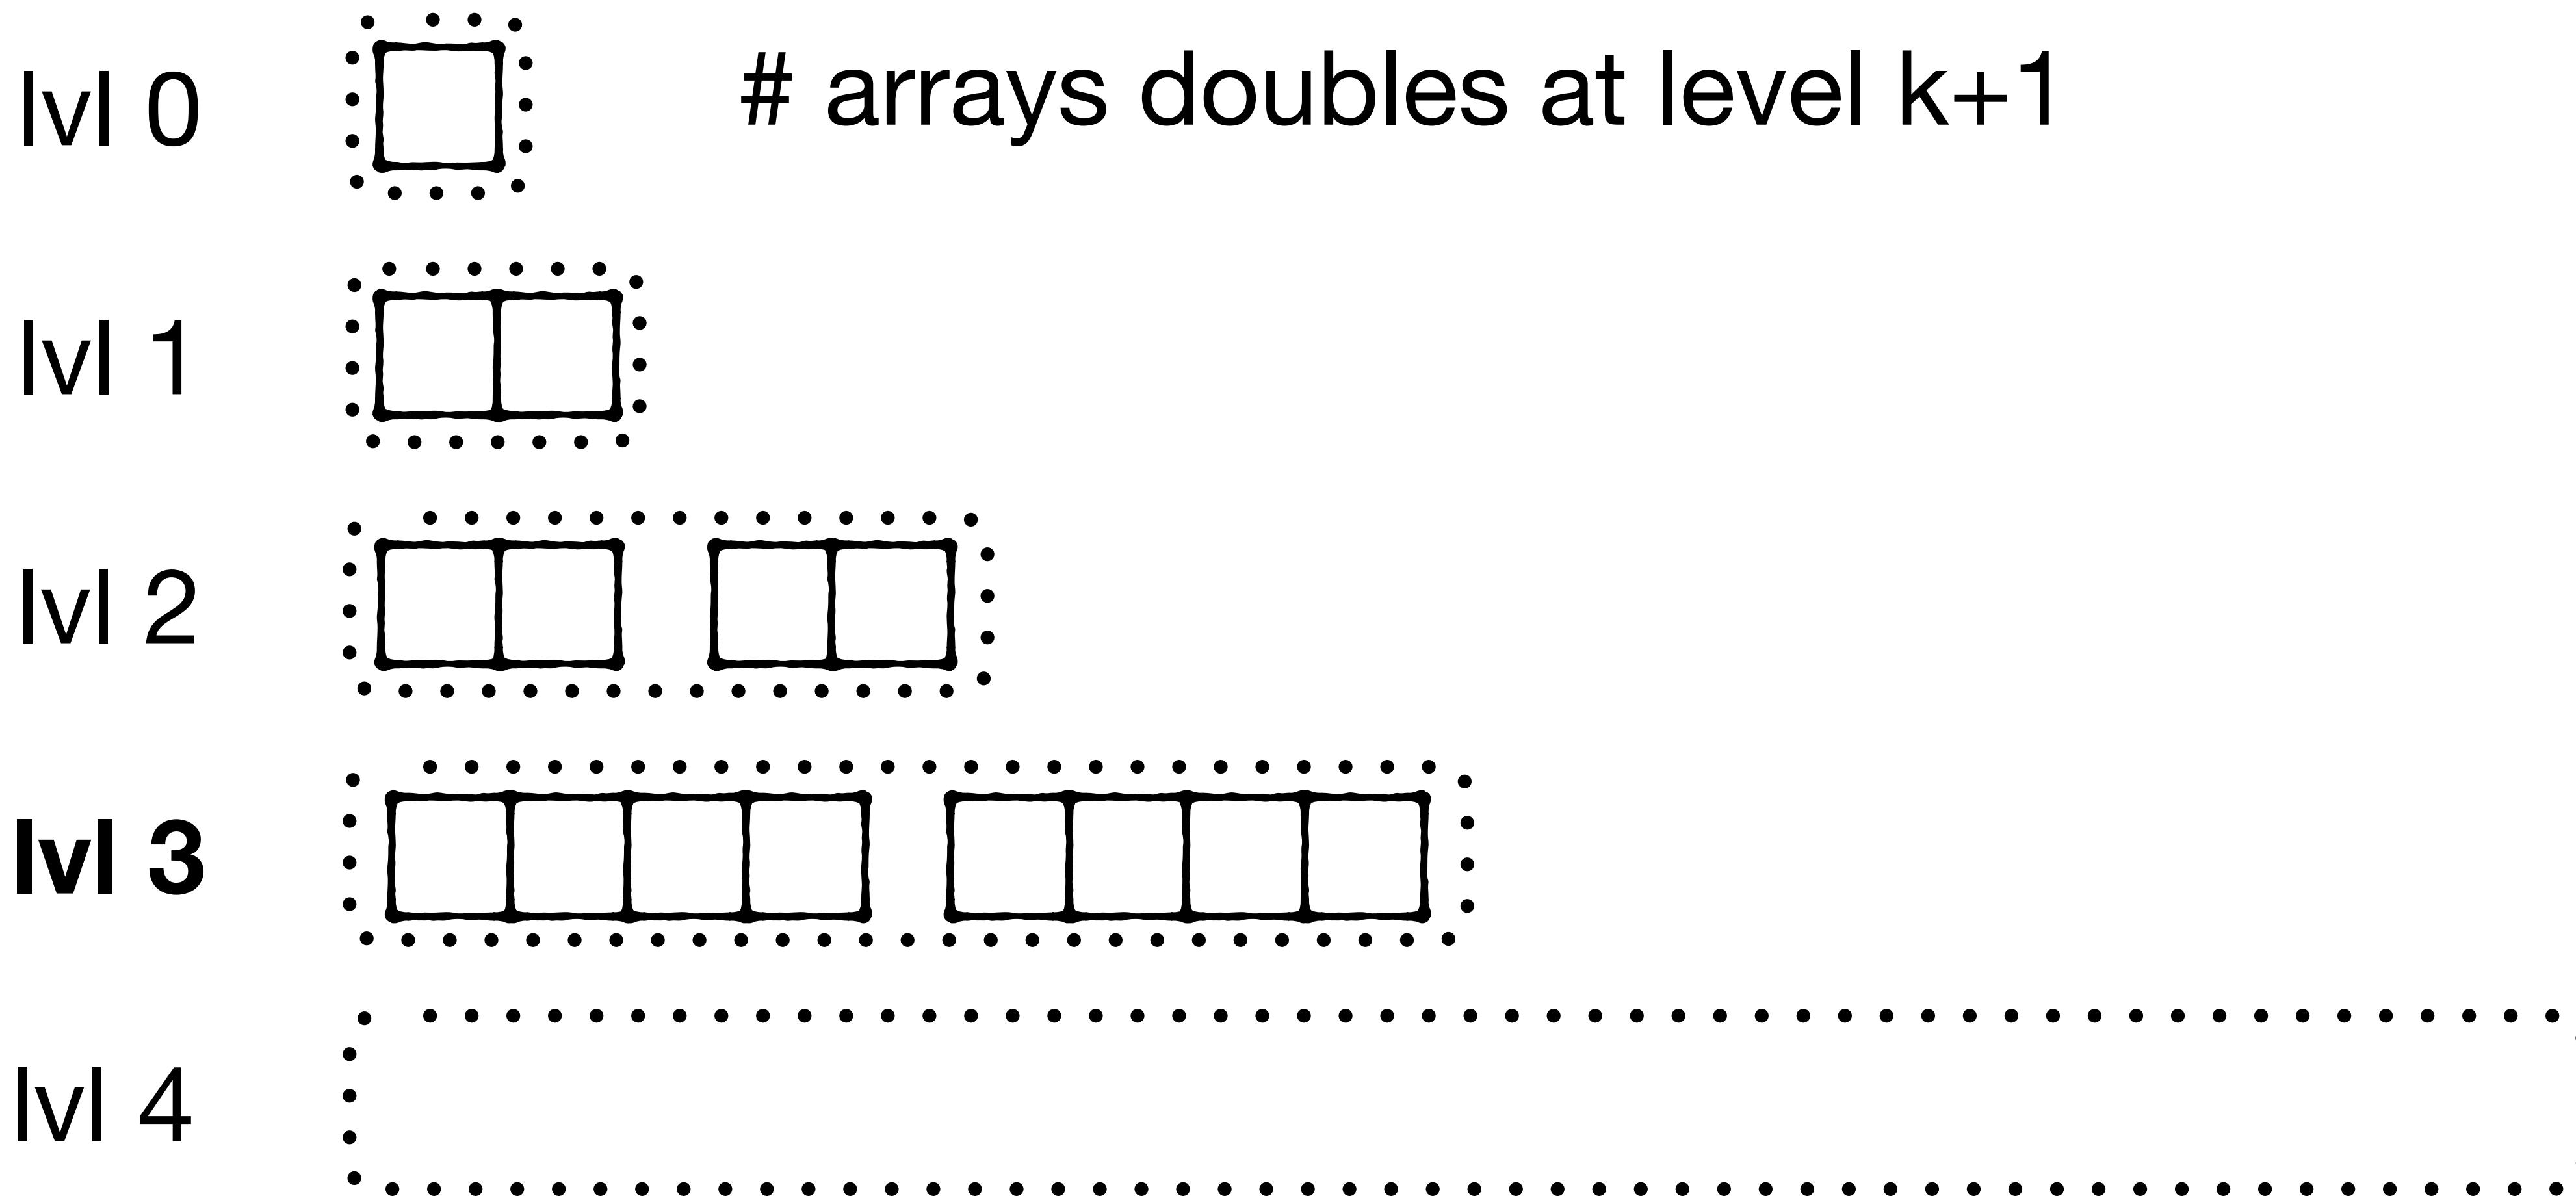
In every pair of subsequent levels k and $k+1$

Size of arrays doubles at level k



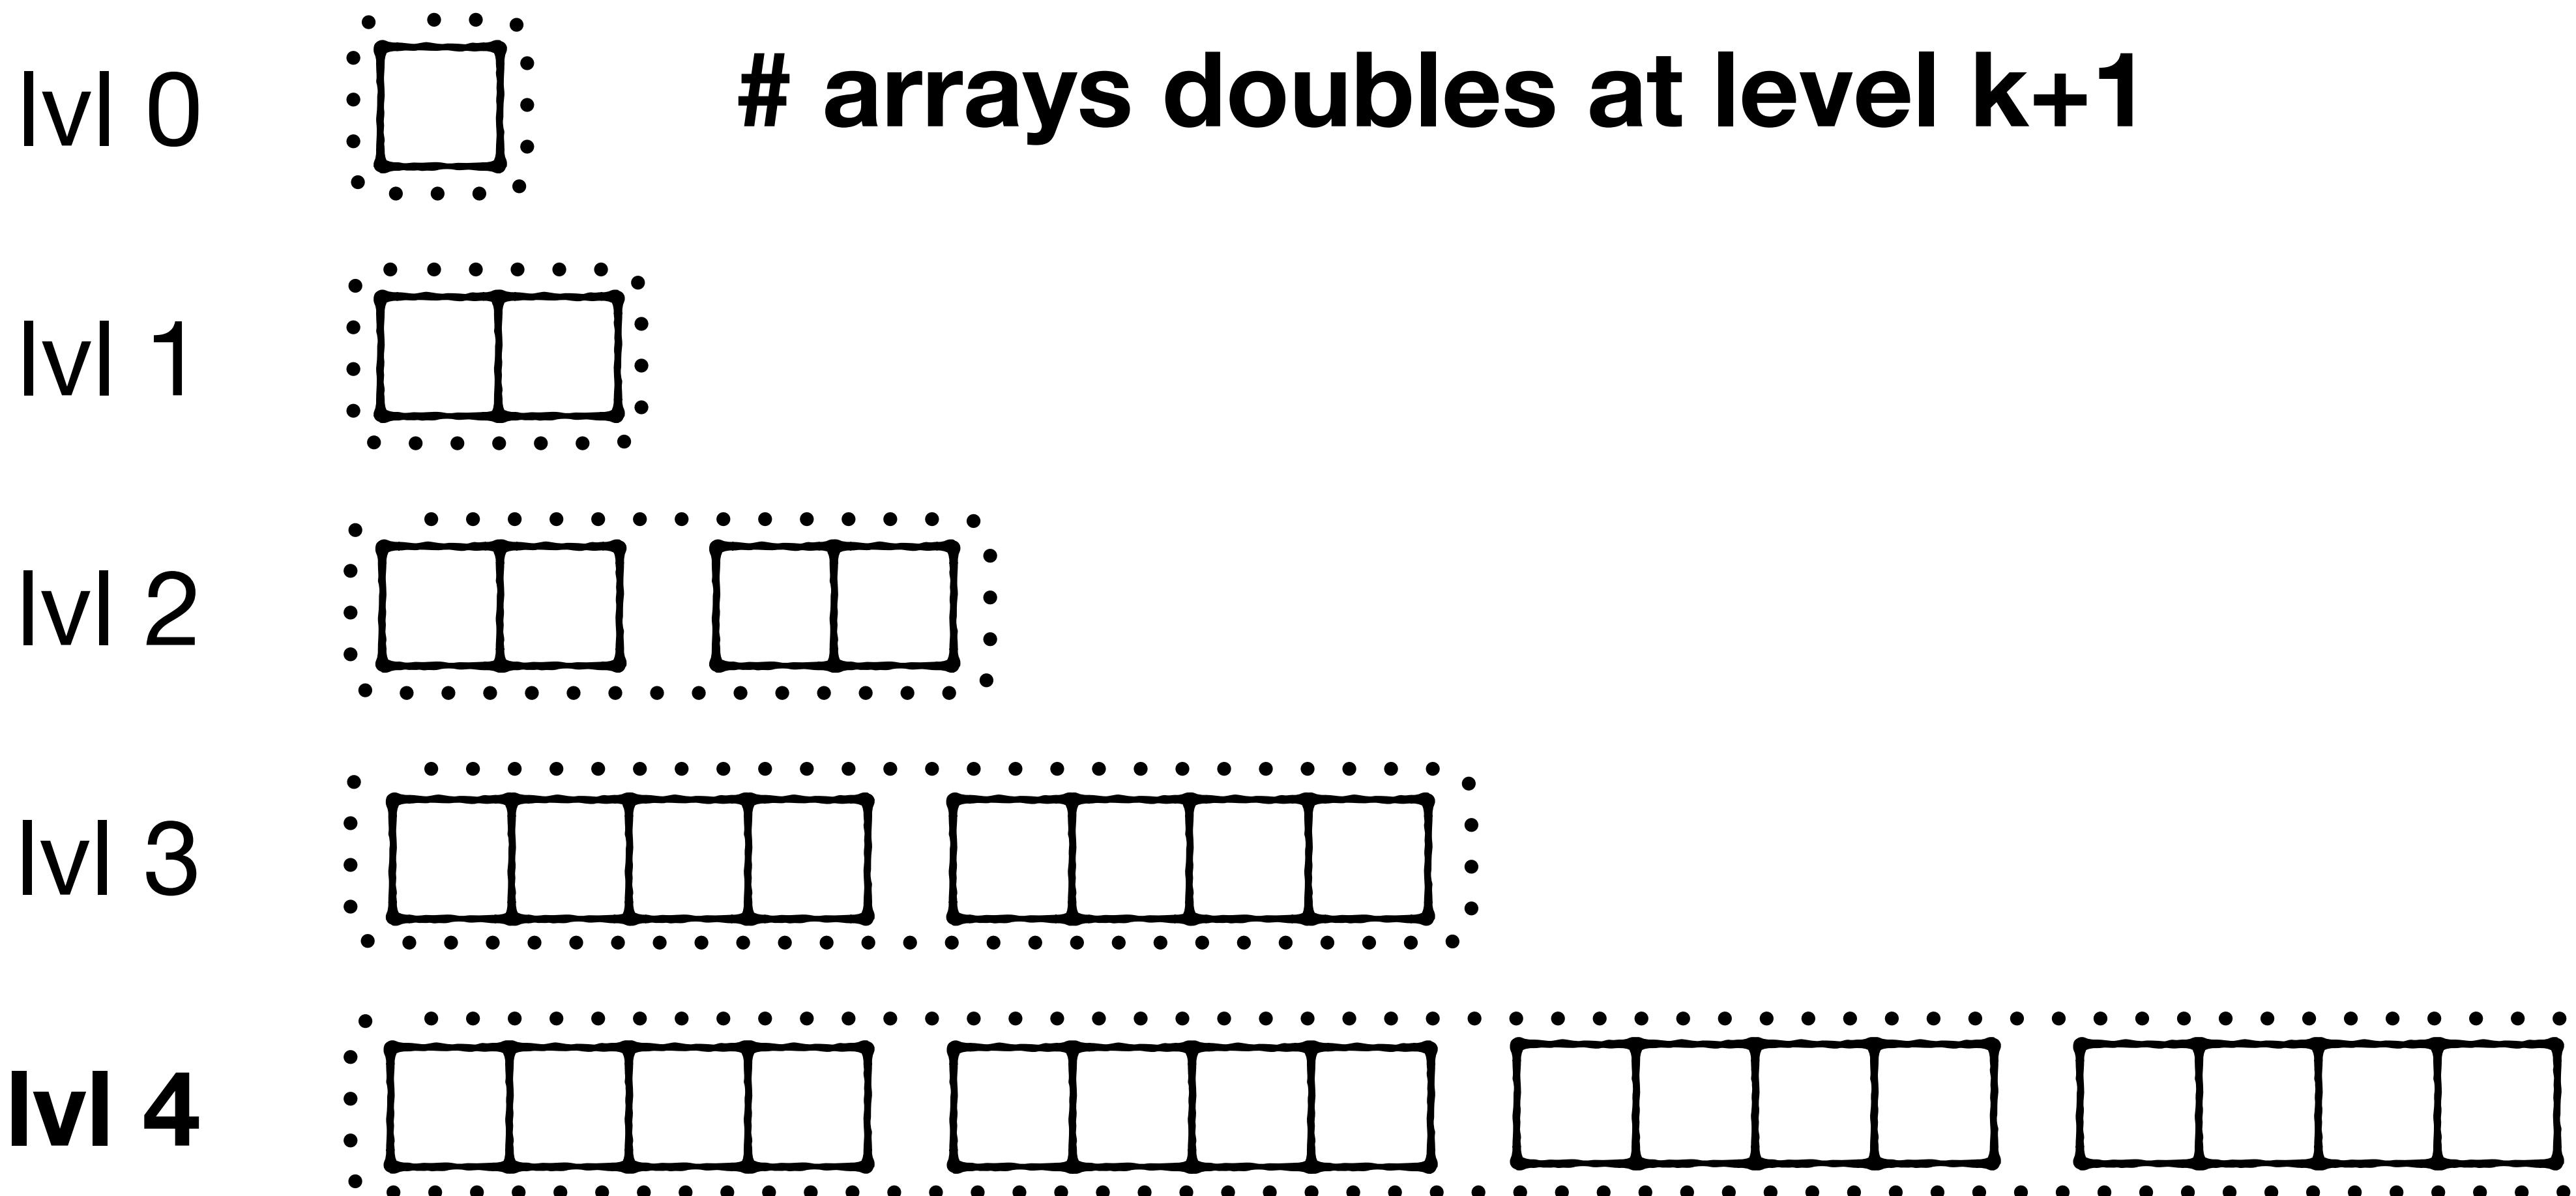
In every pair of subsequent levels k and $k+1$

Size of arrays doubles at level k

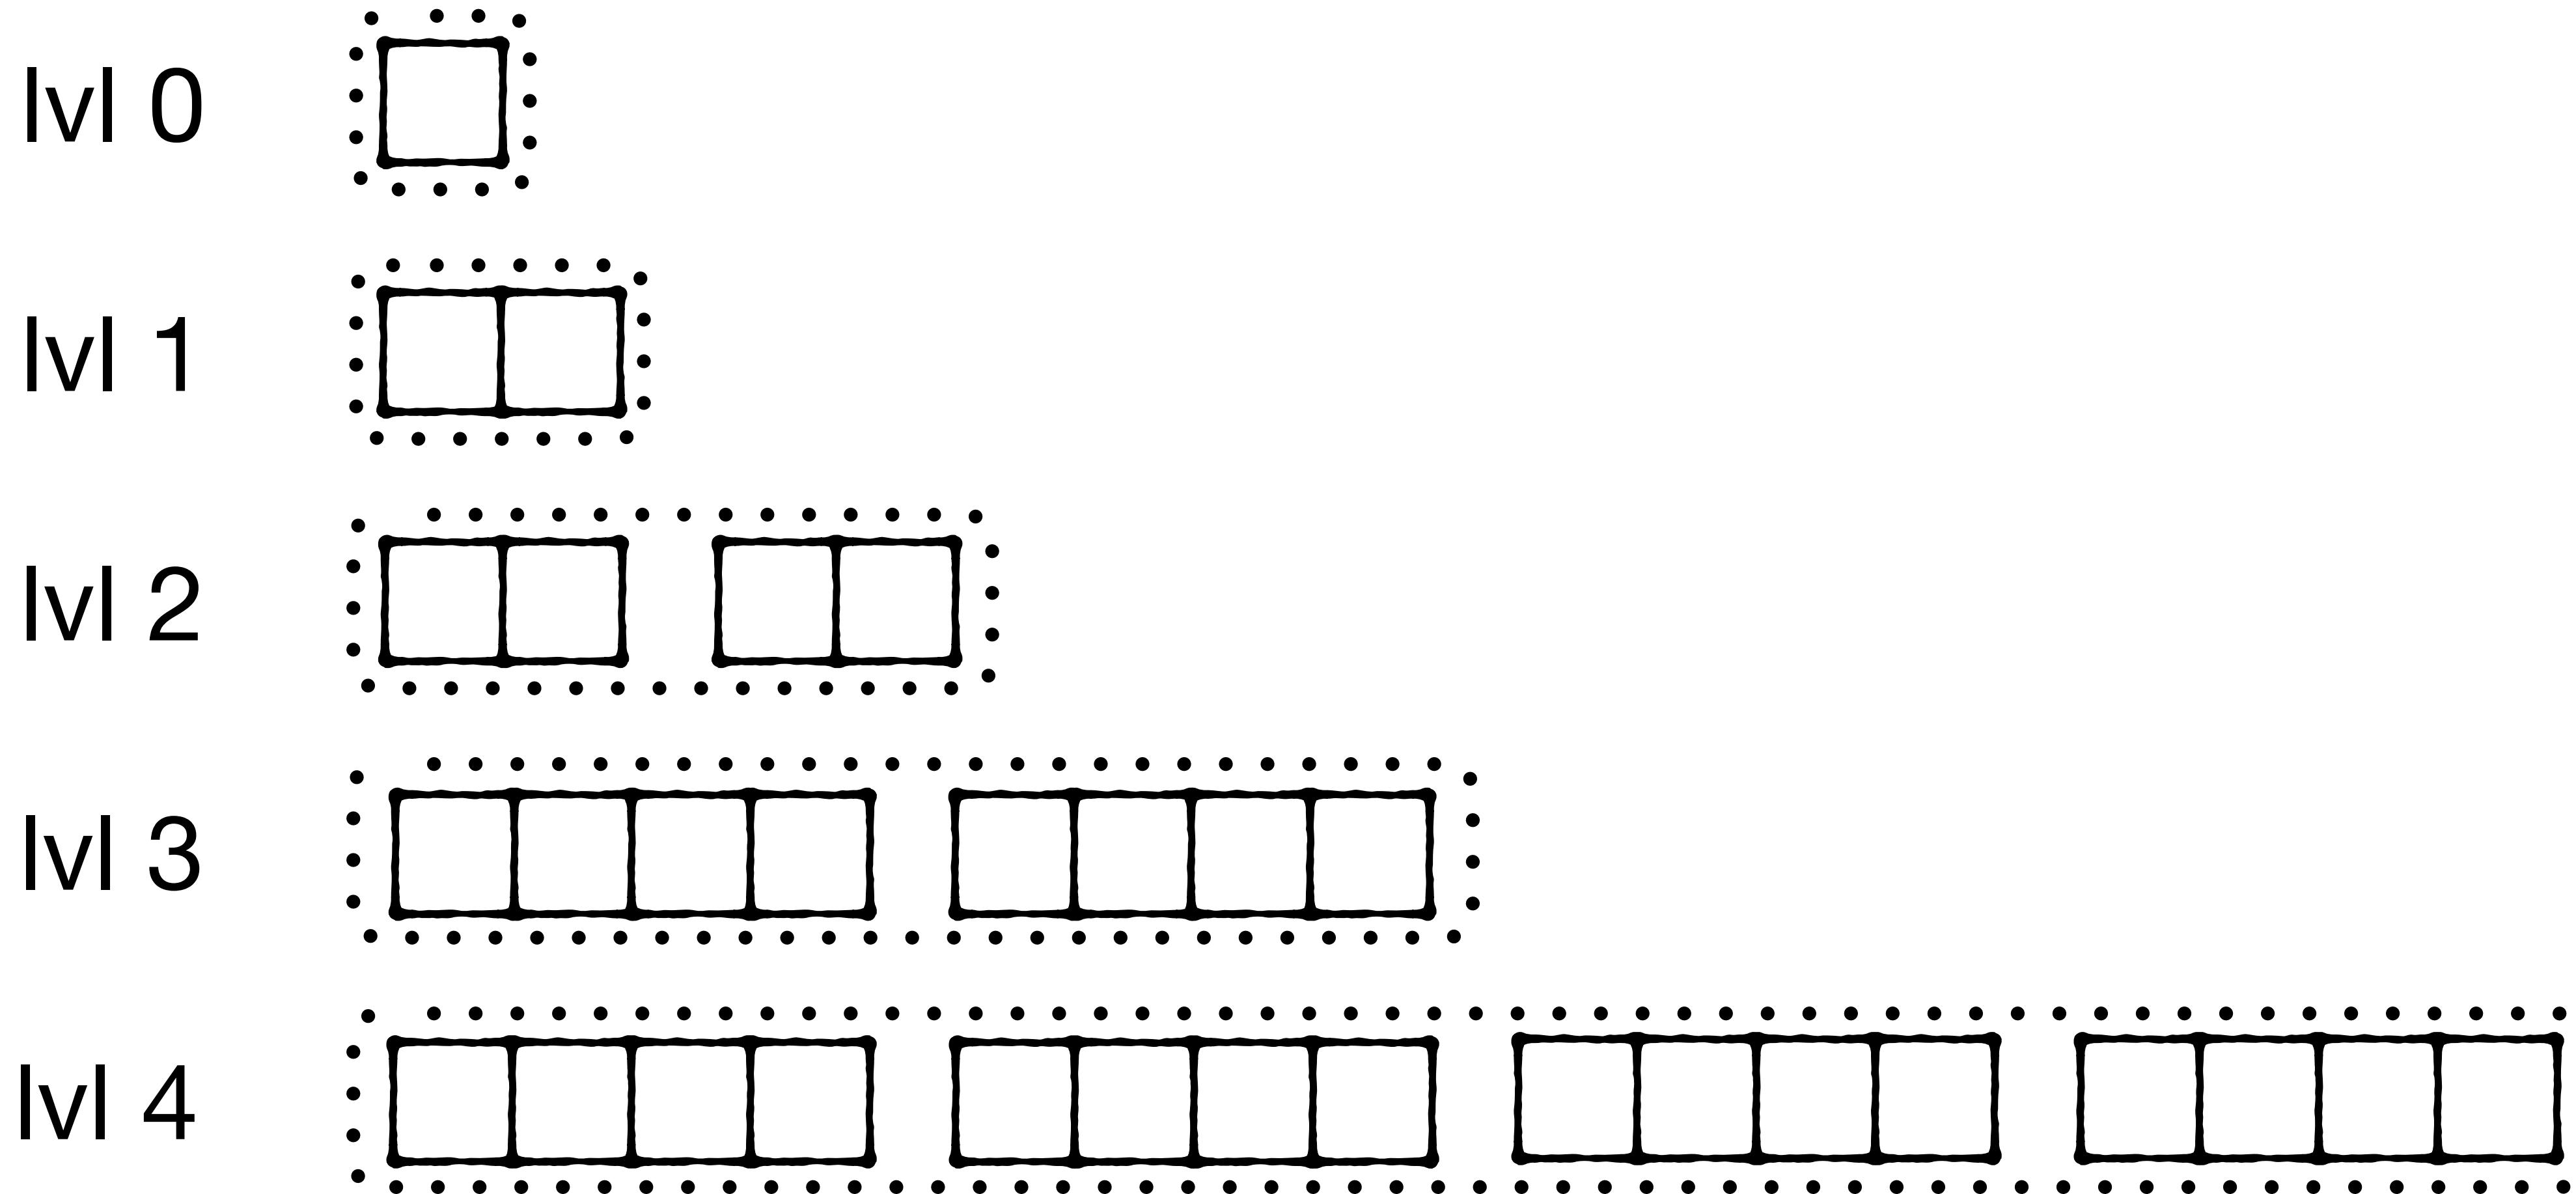


In every pair of subsequent levels k and $k+1$

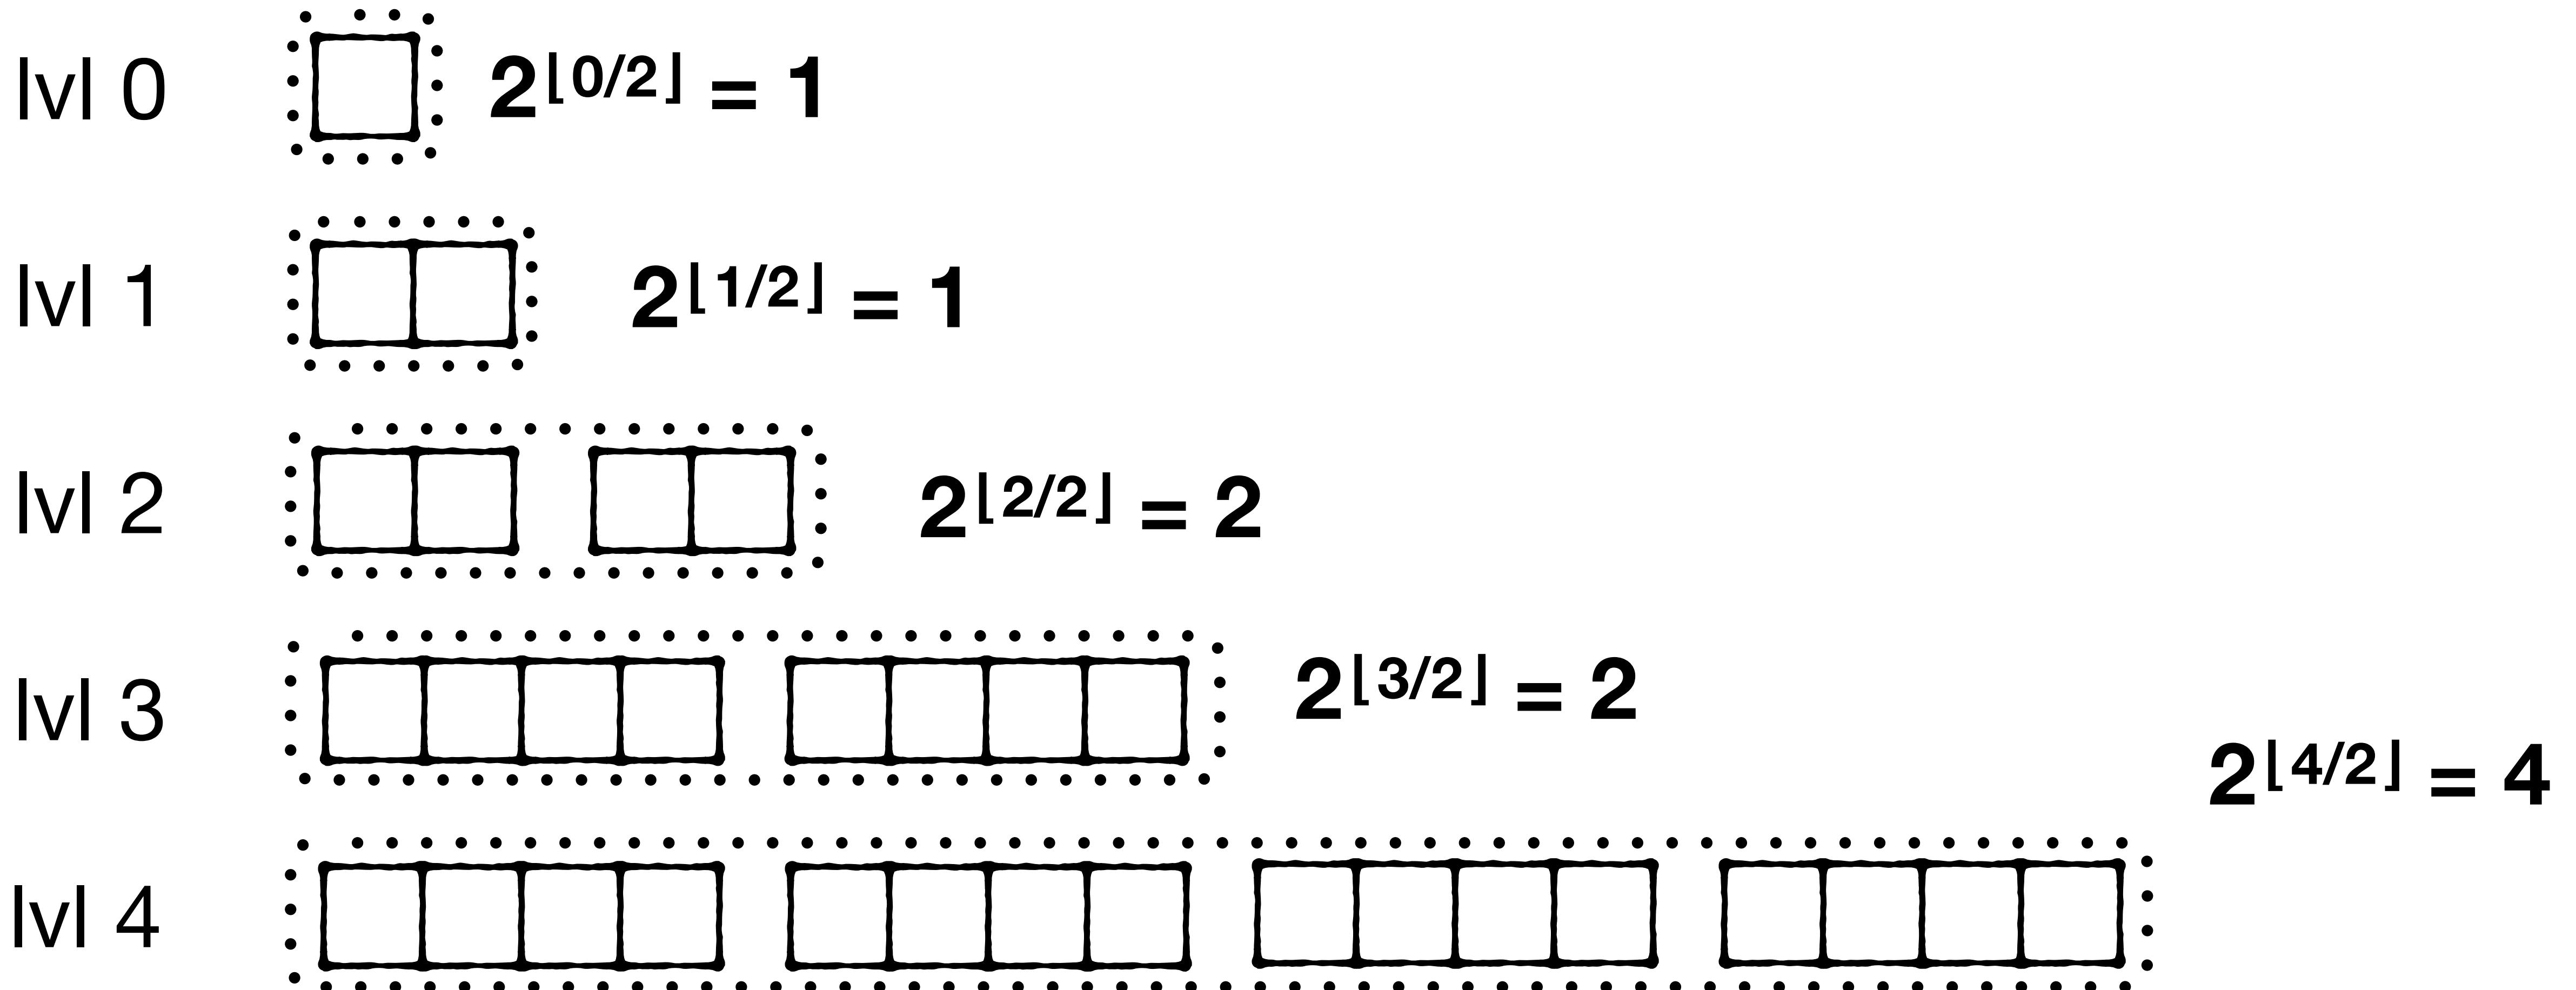
Size of arrays doubles at level k



lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

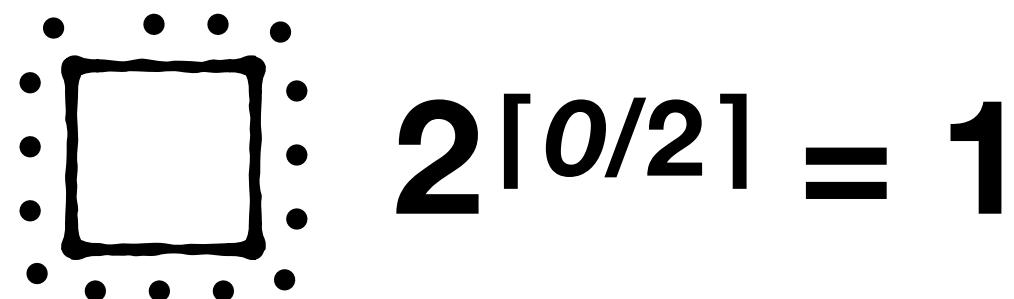


lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots



lvl i contains $2^{\lceil K/2 \rceil}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

lvl 0



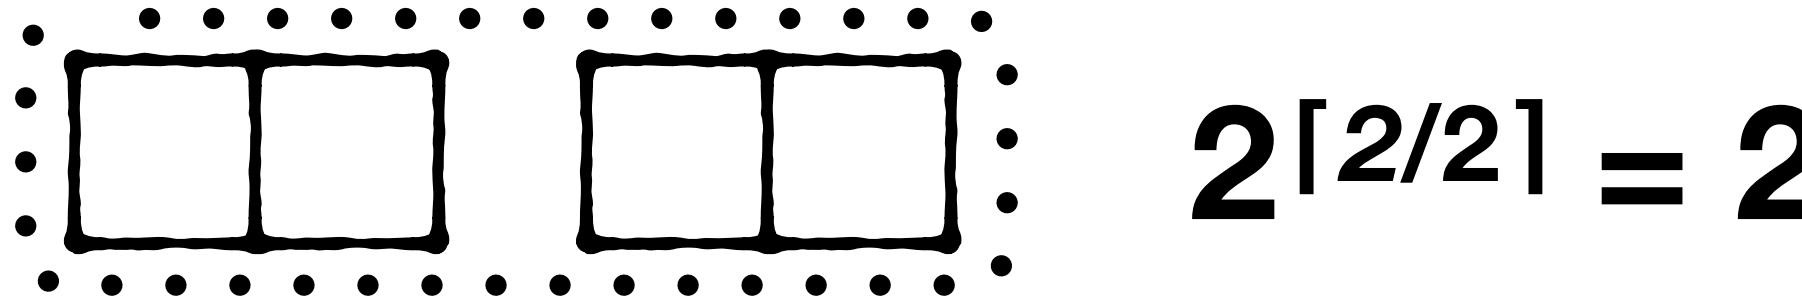
$$2^{\lceil 0/2 \rceil} = 1$$

lvl 1



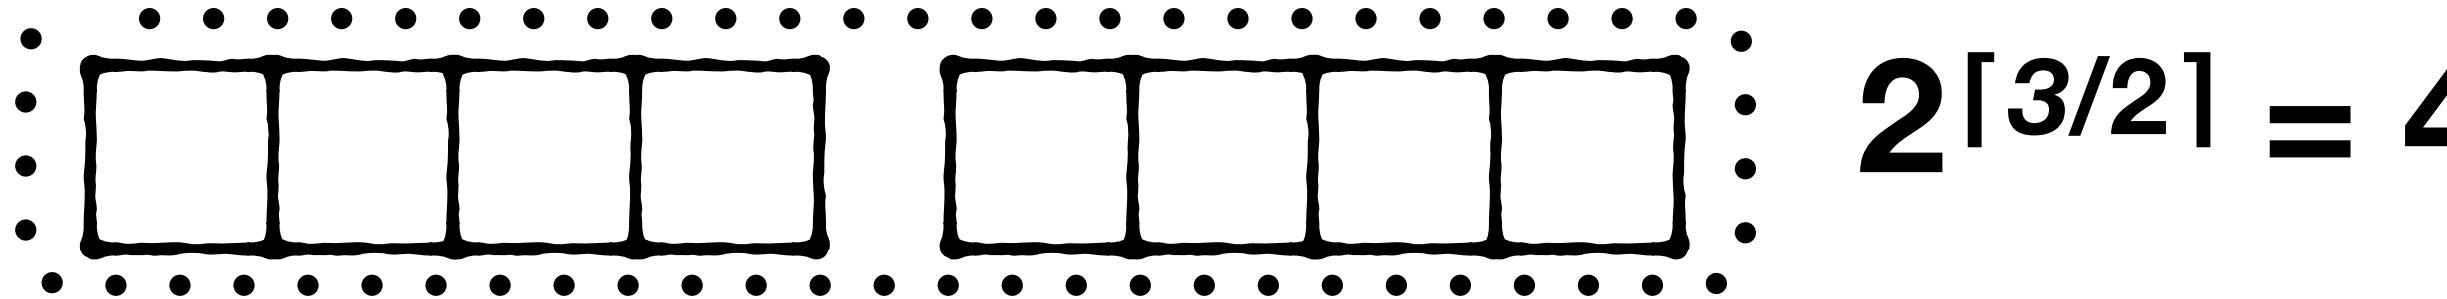
$$2^{\lceil 1/2 \rceil} = 2$$

lvl 2



$$2^{\lceil 2/2 \rceil} = 2$$

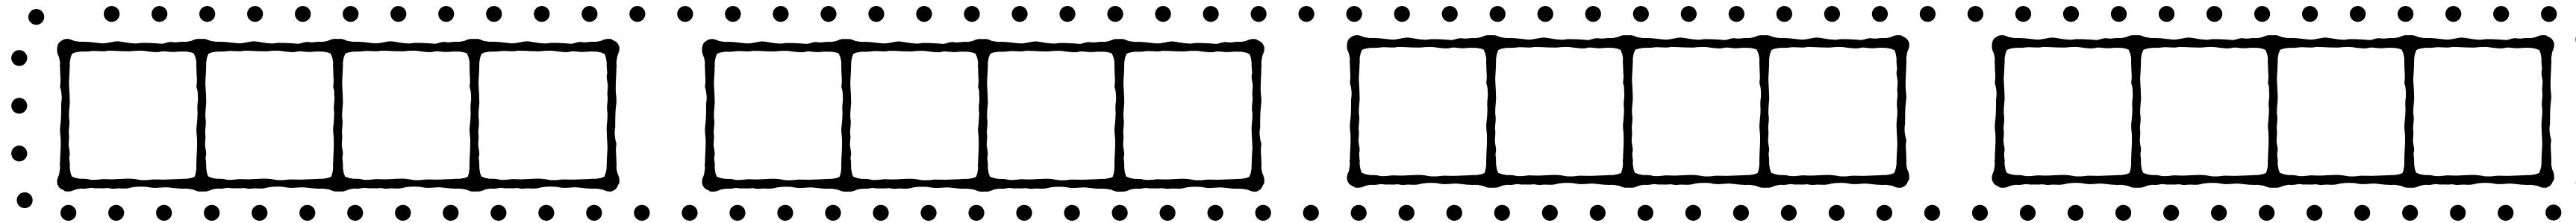
lvl 3



$$2^{\lceil 3/2 \rceil} = 4$$

$$2^{\lceil 4/2 \rceil} = 4$$

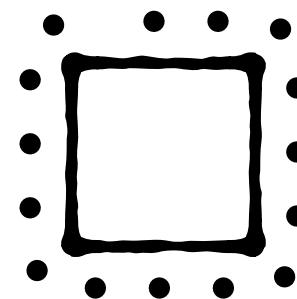
lvl 4



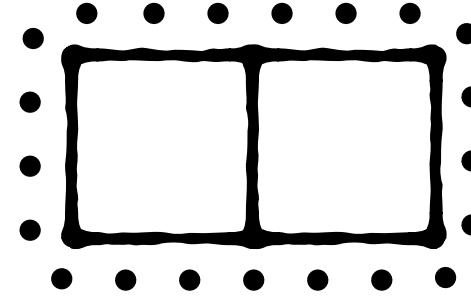
lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

levels?

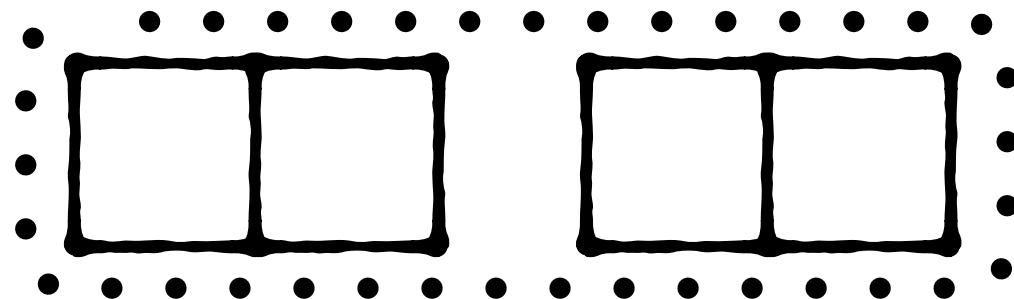
lvl 0



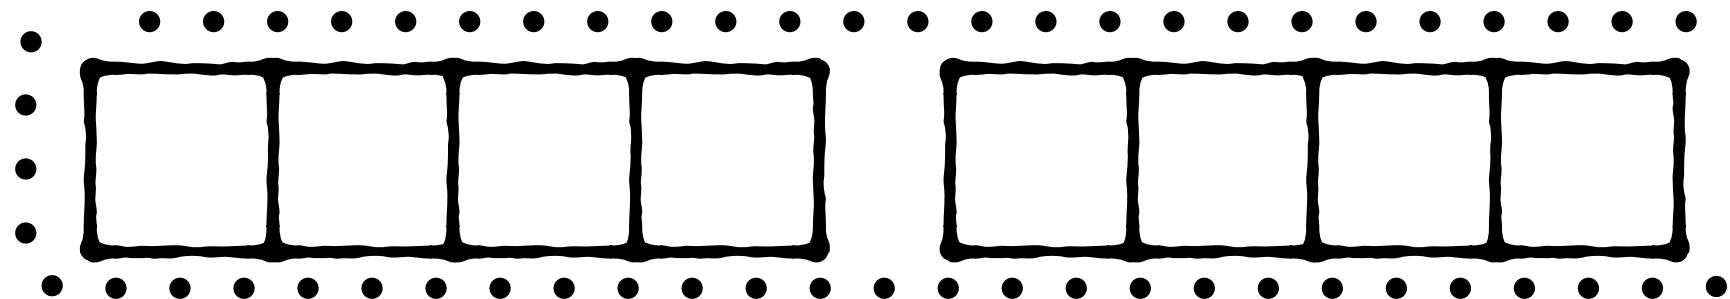
lvl 1



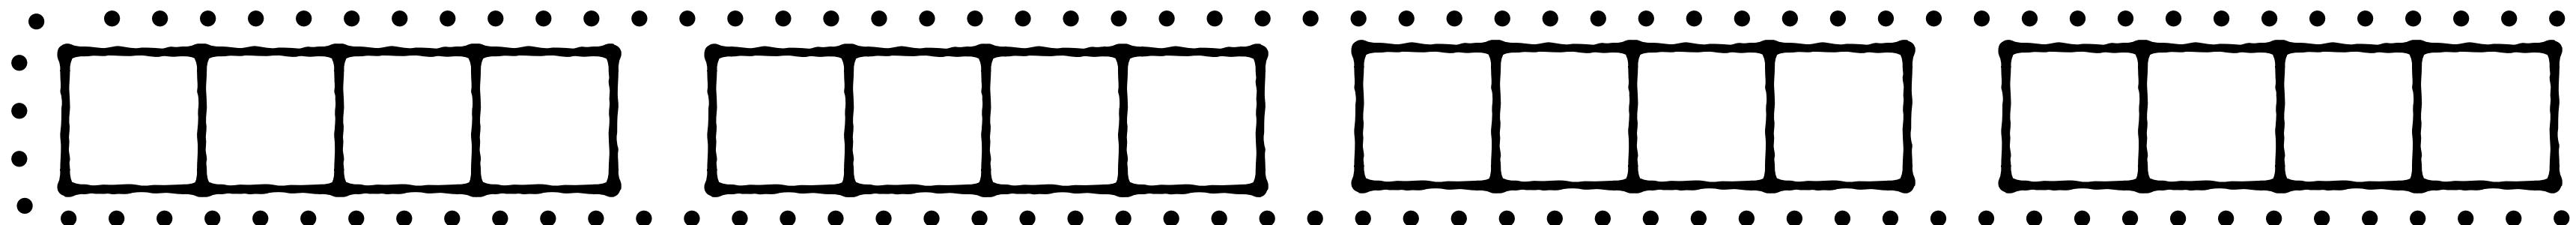
lvl 2



lvl 3

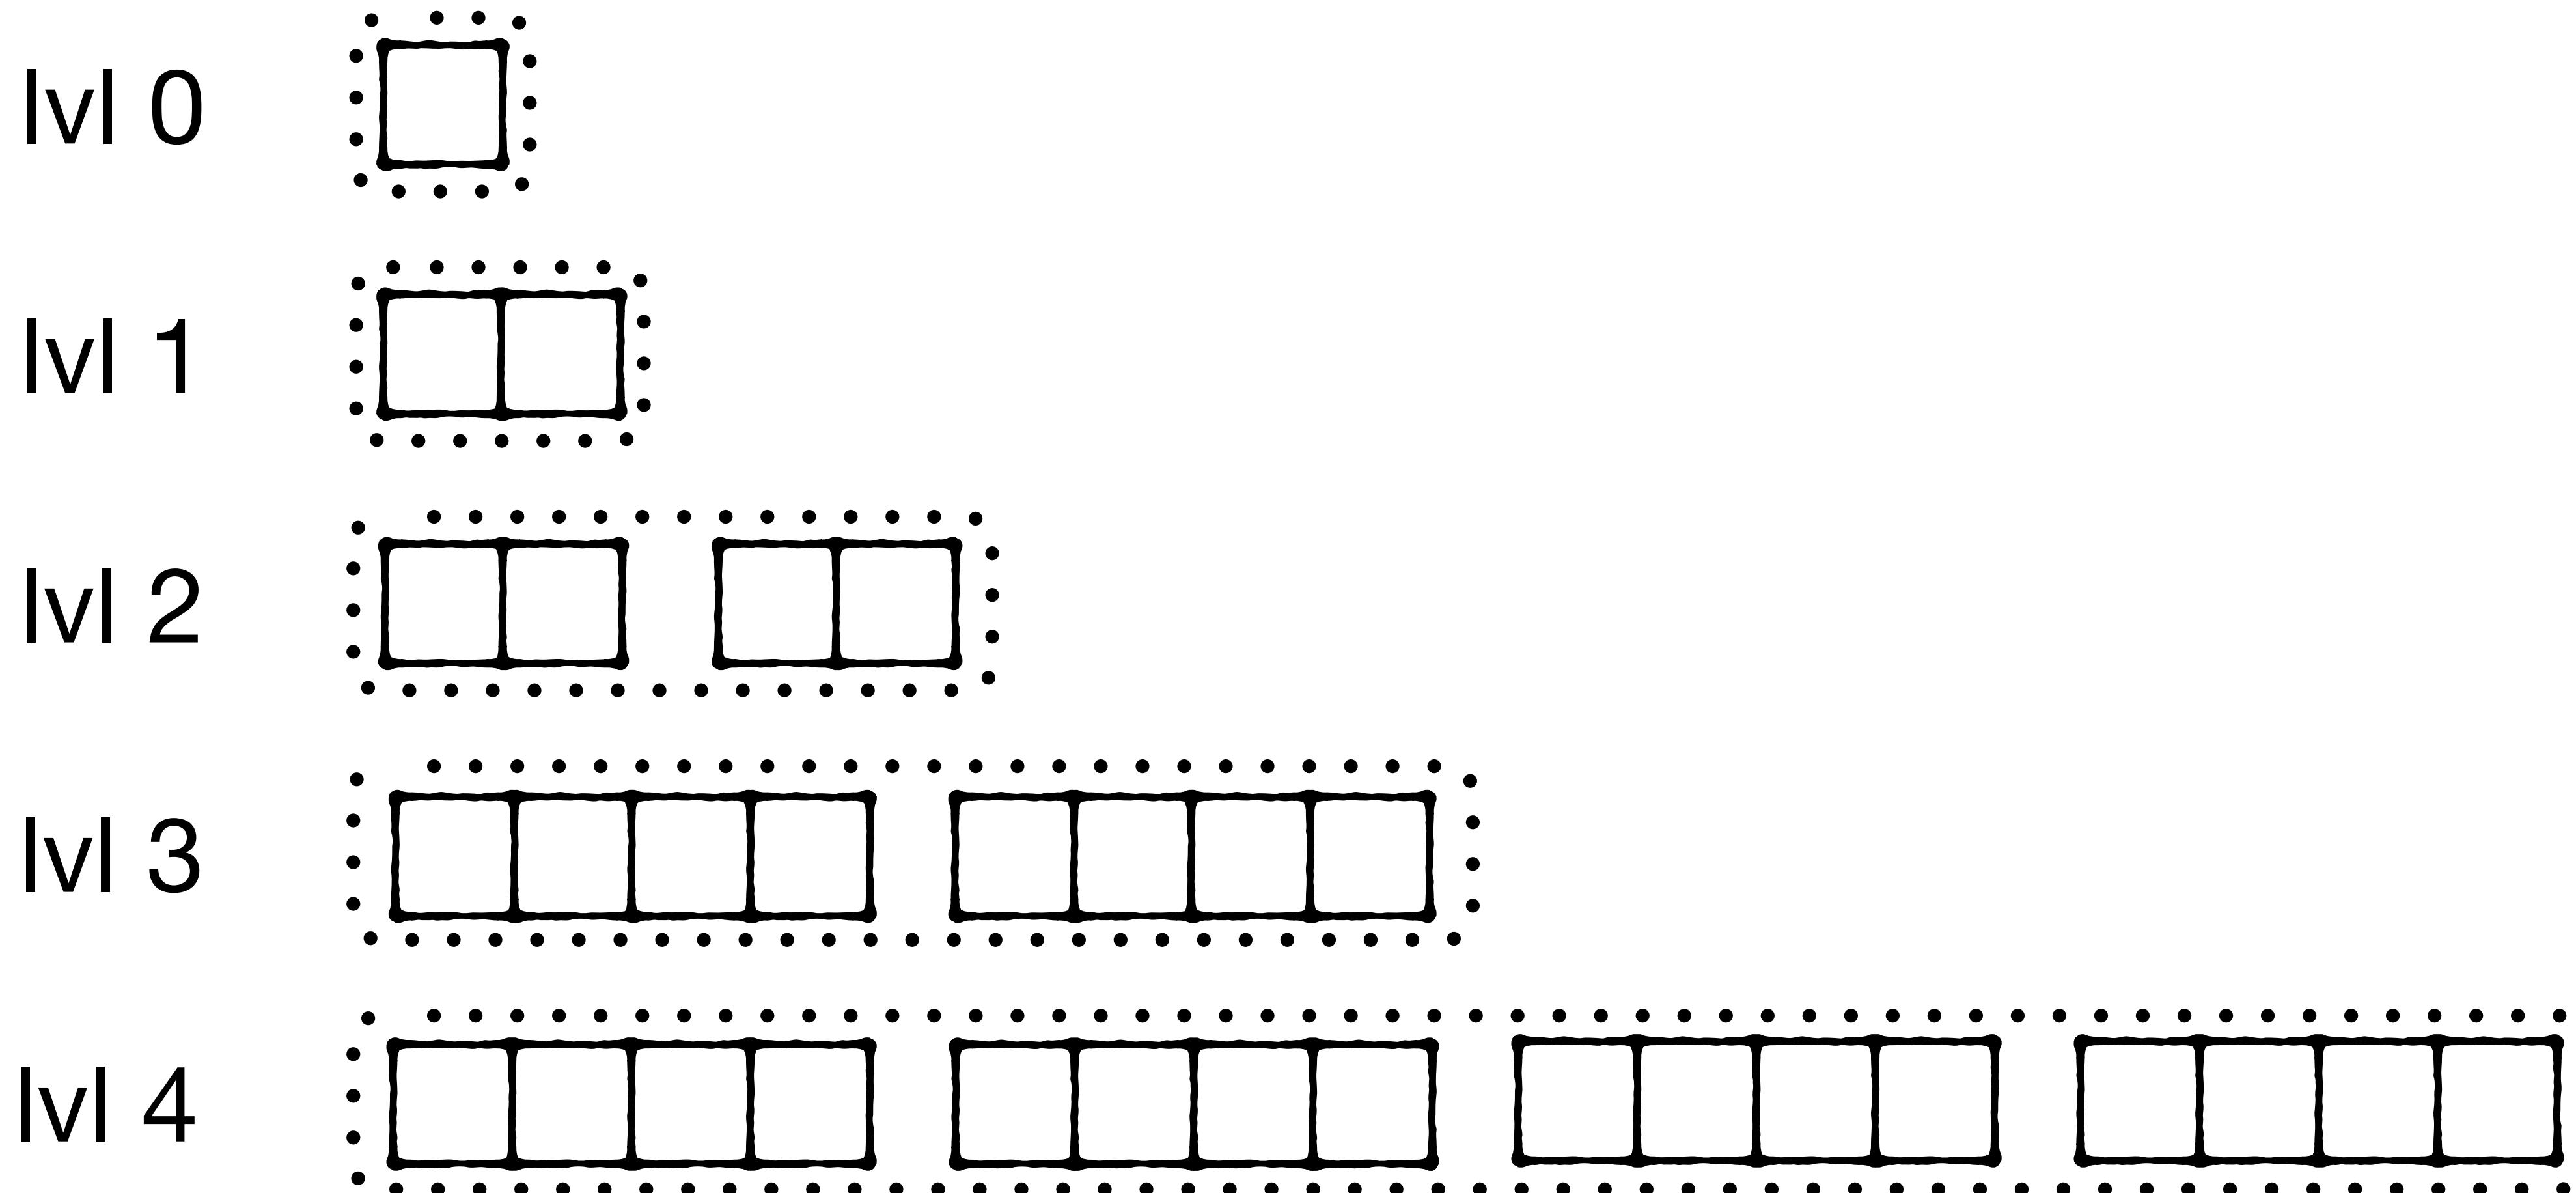


lvl 4



lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

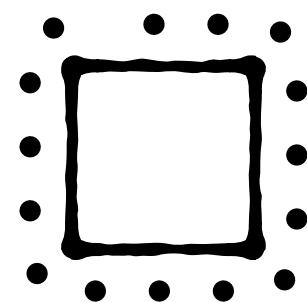
levels: $\log_2 N$



lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

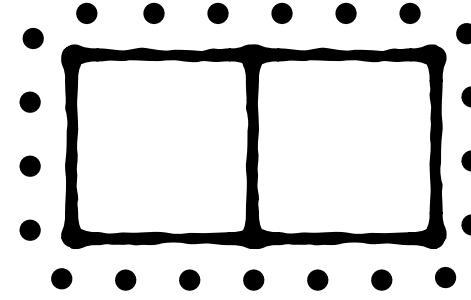
levels: $\log_2 N$

lvl 0

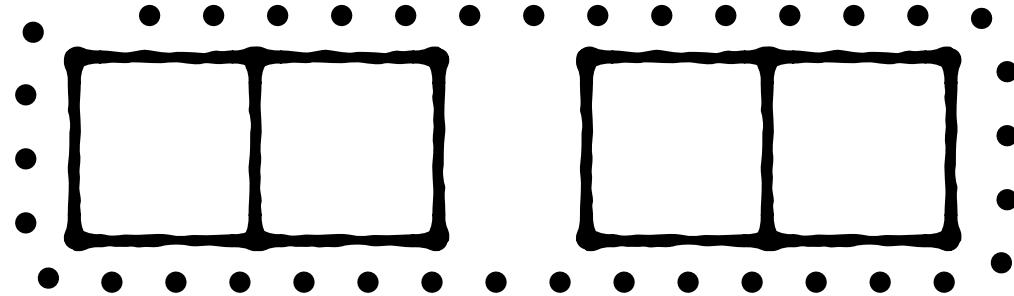


data blocks?

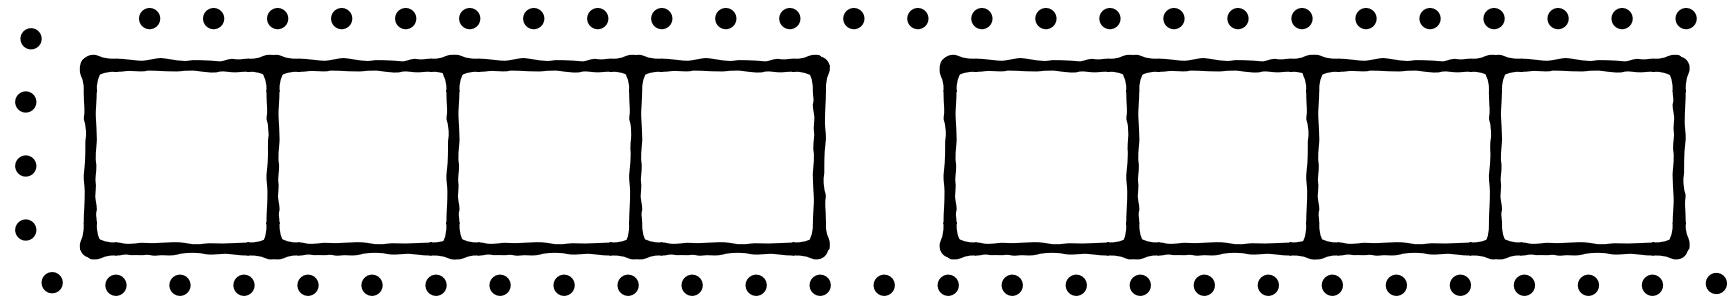
lvl 1



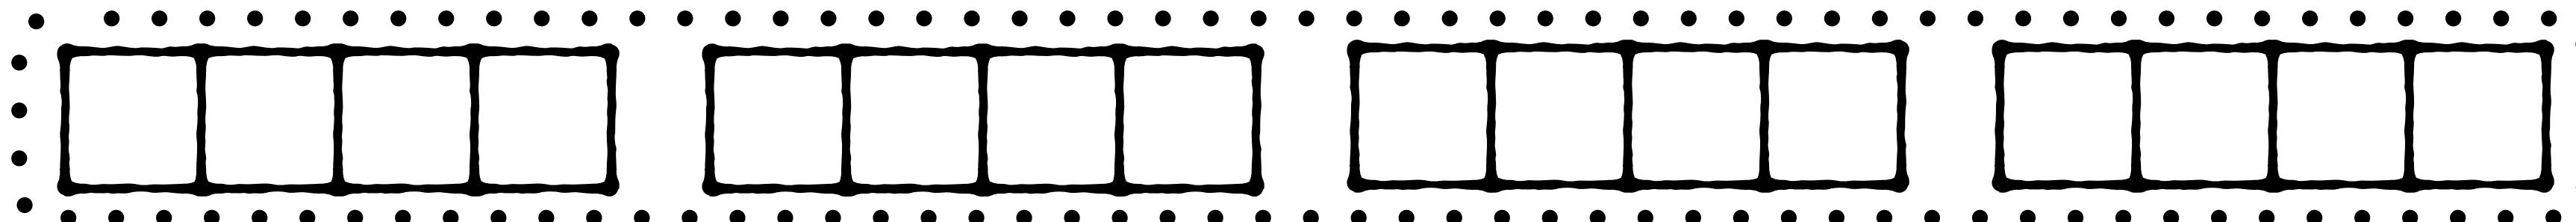
lvl 2



lvl 3

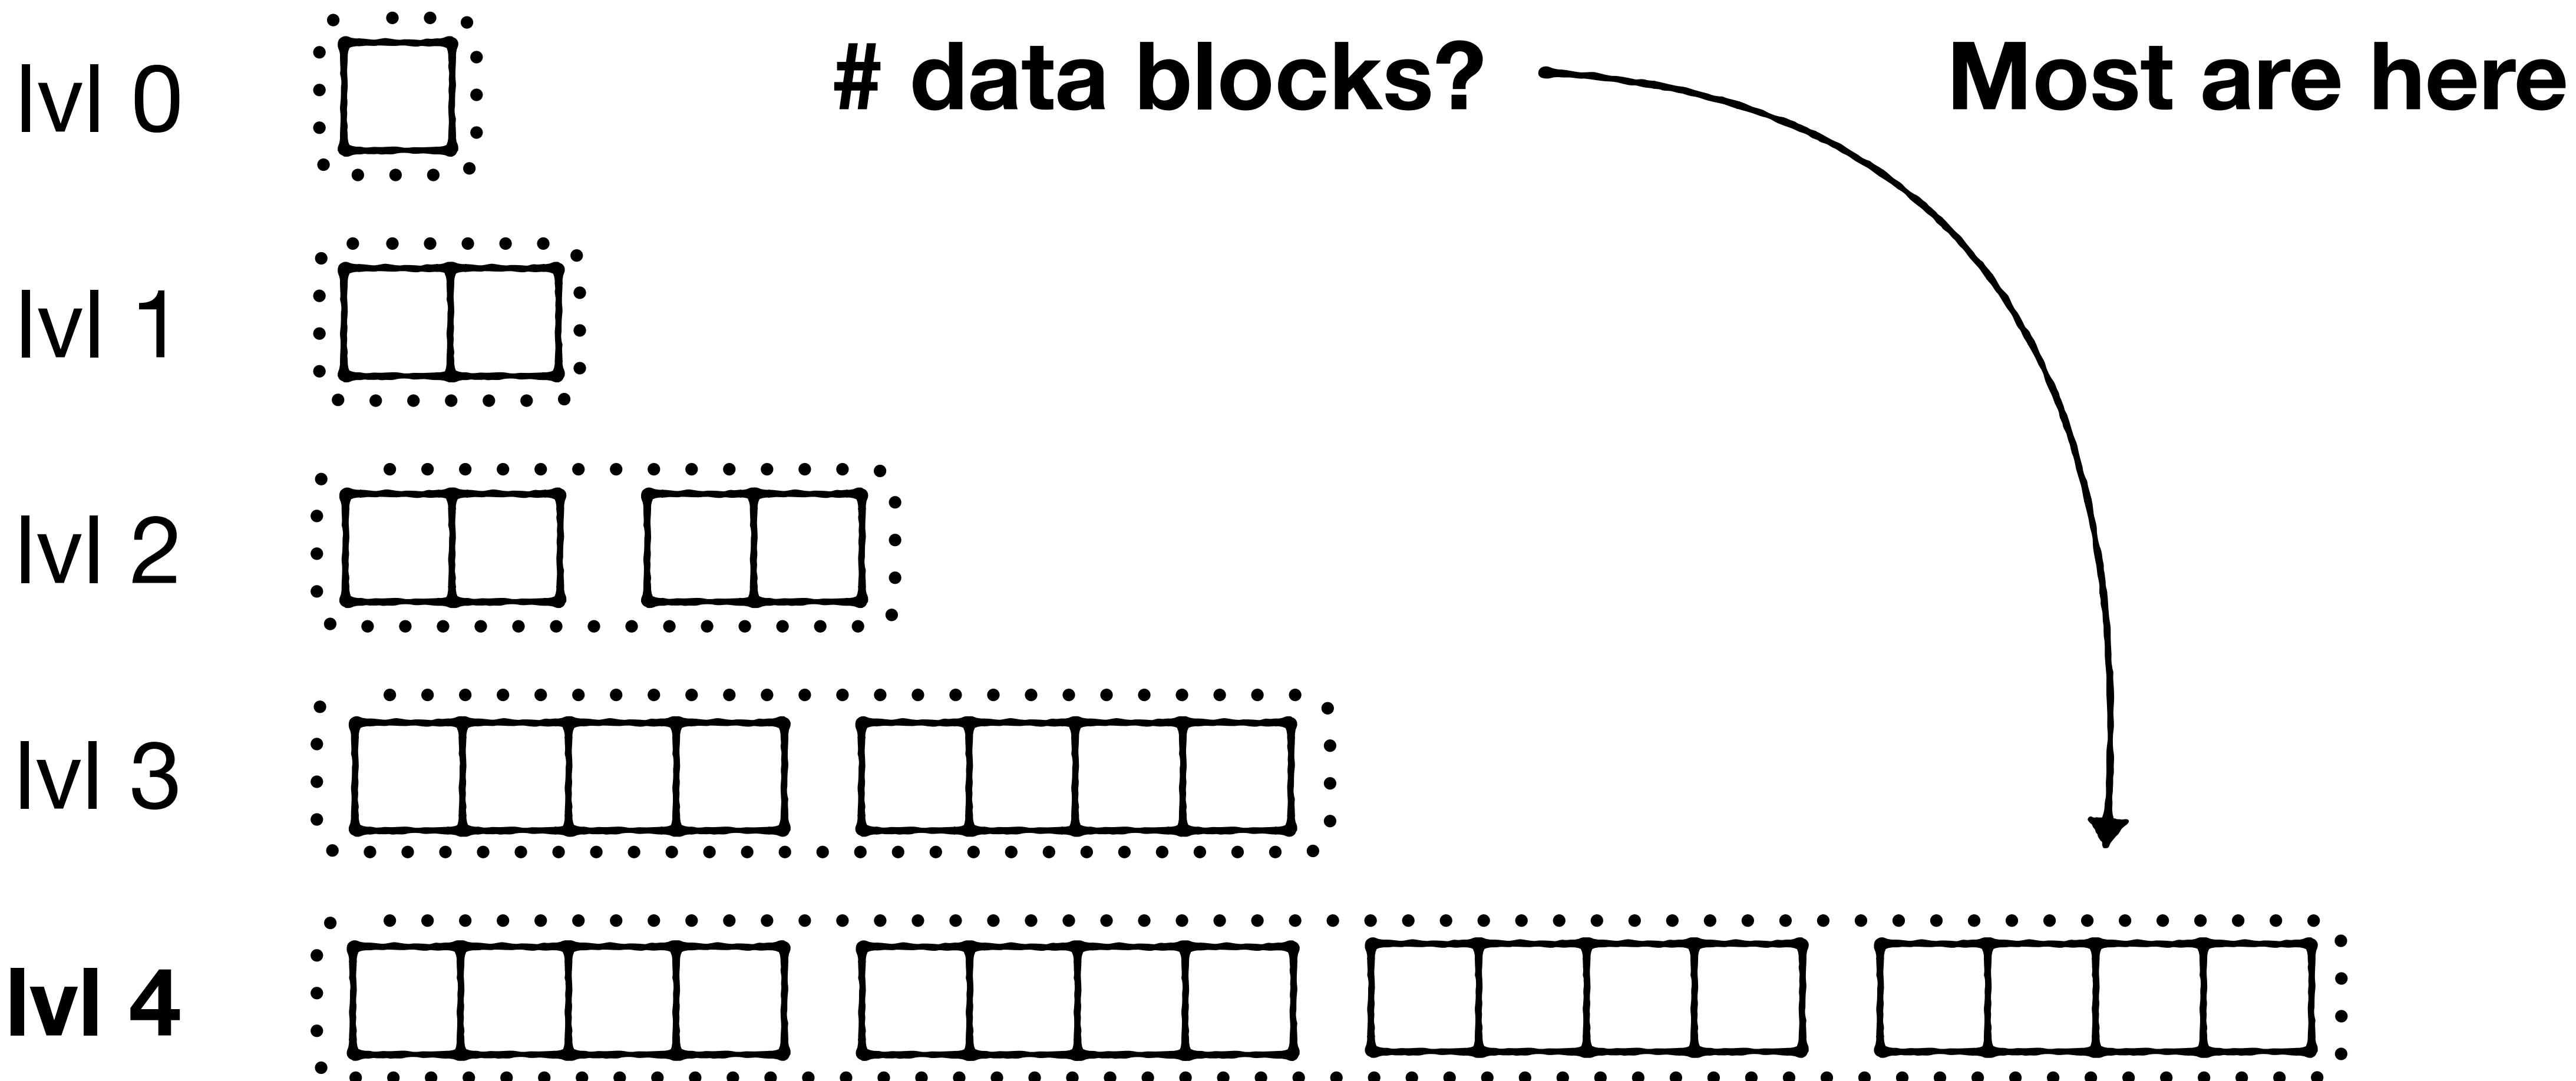


lvl 4



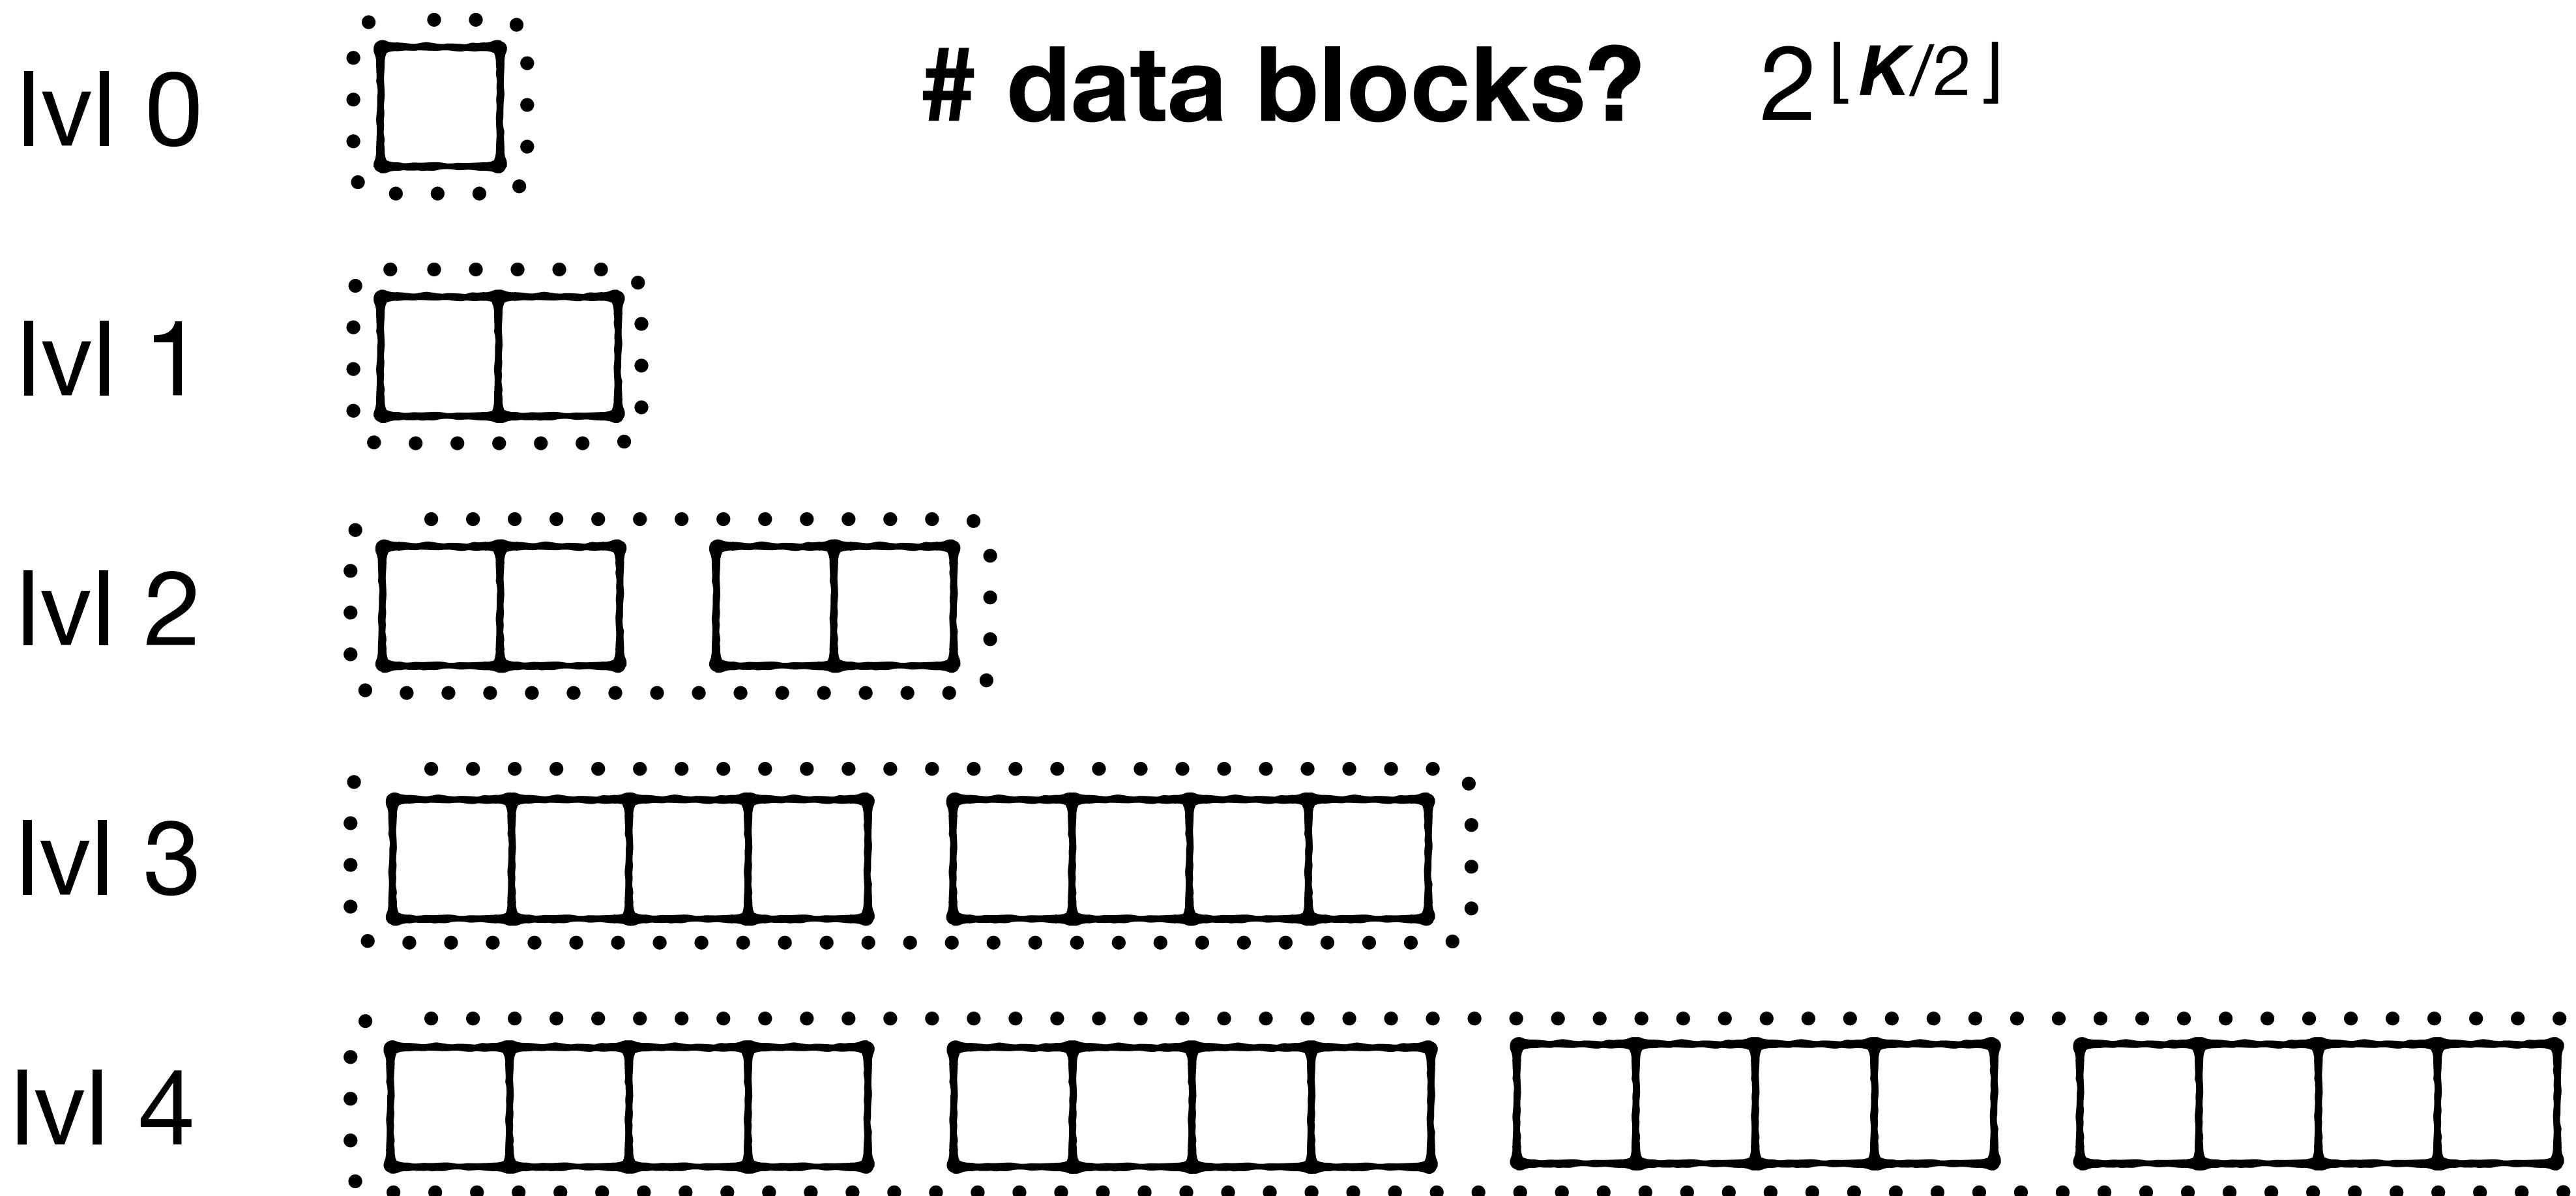
lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

levels: $\log_2 N$



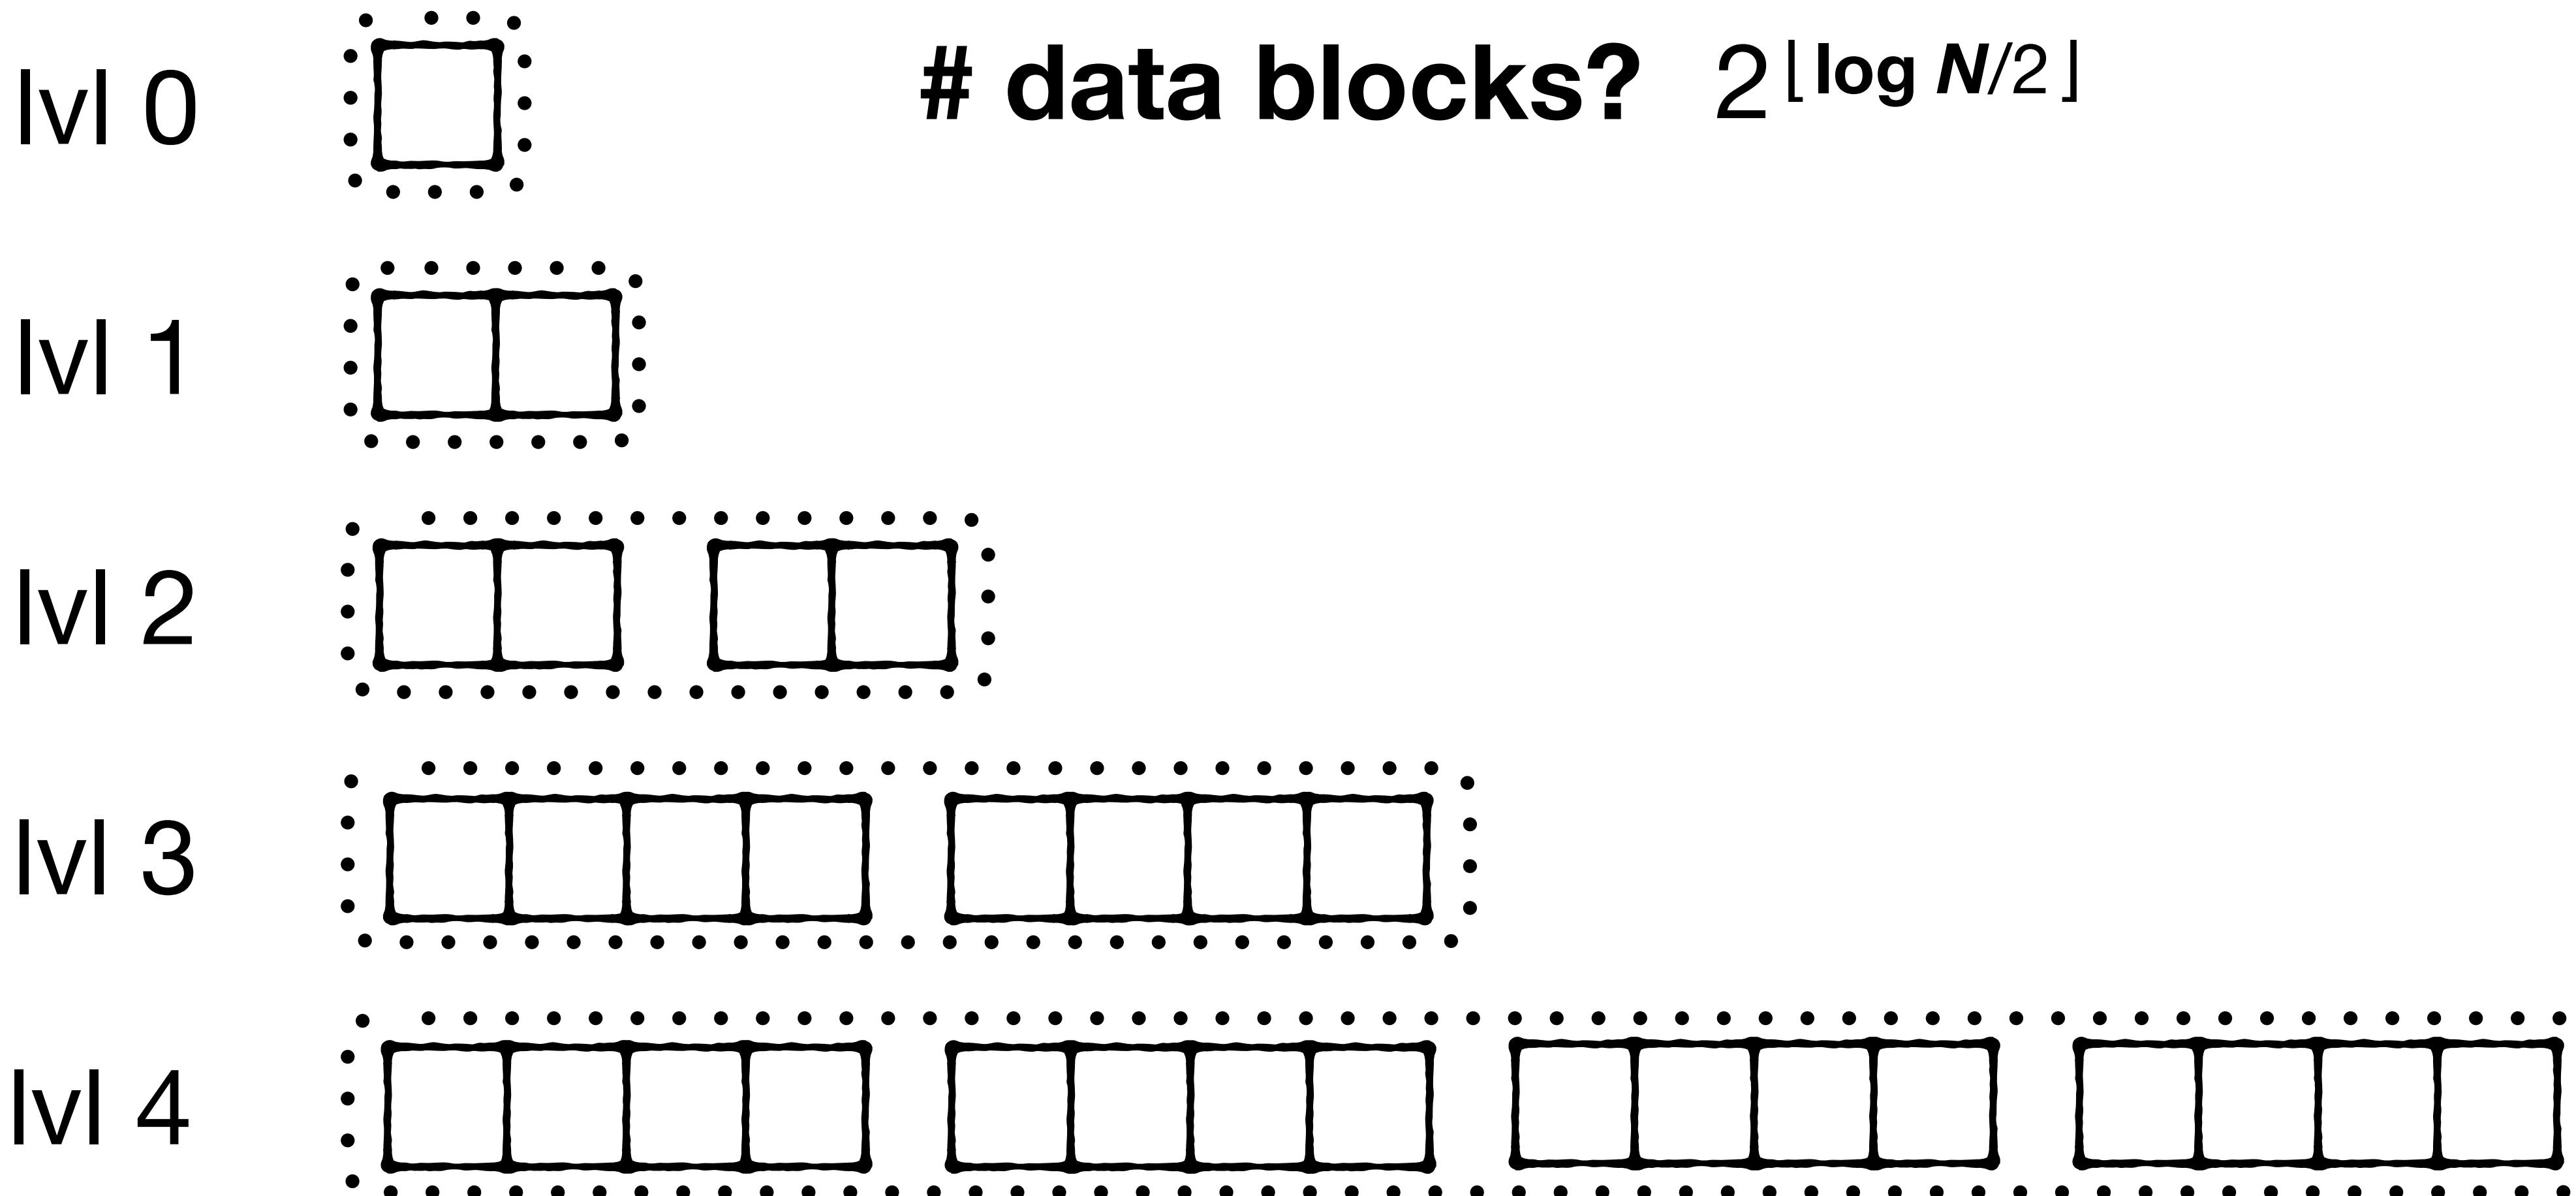
lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

levels: $\log_2 N$ ↗



lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

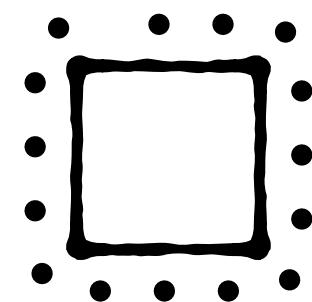
levels: $\log_2 N$ 



lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

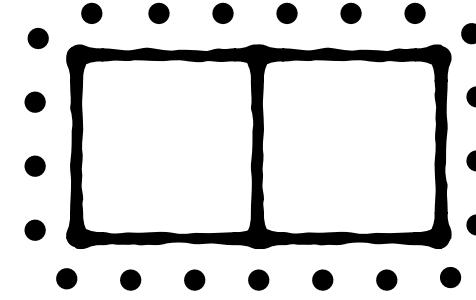
levels: $\log_2 N$

lvl 0

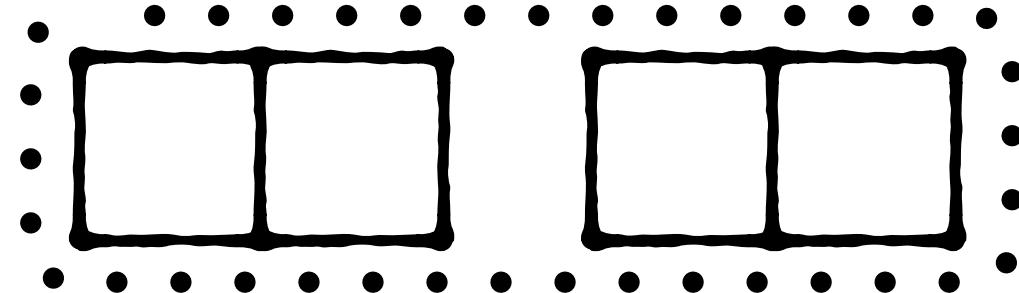


data blocks? $O(\sqrt{N})$

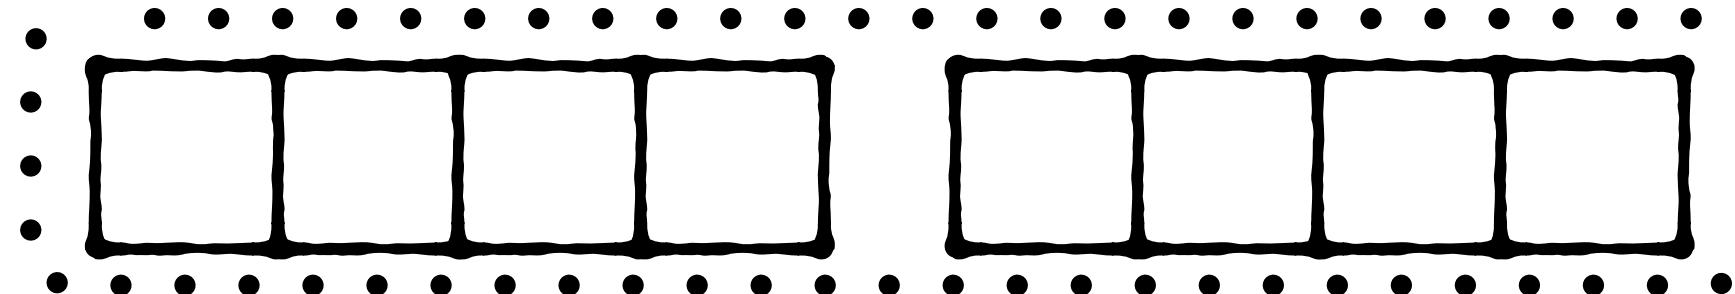
lvl 1



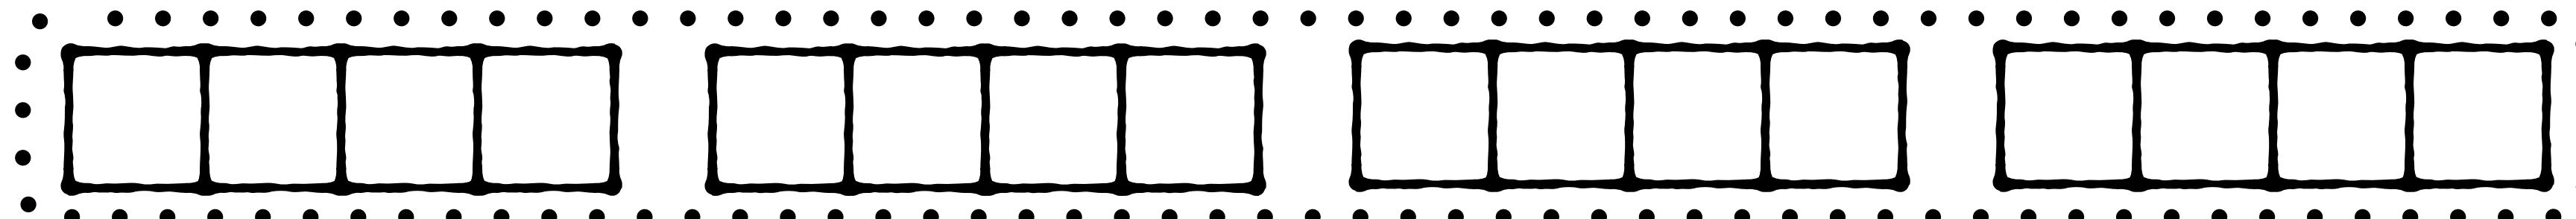
lvl 2



lvl 3



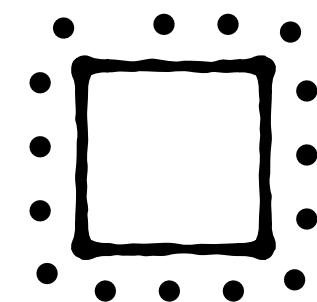
lvl 4



lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

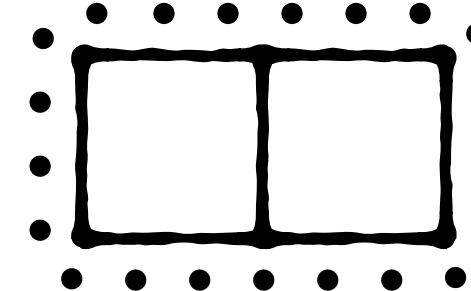
levels: $\log_2 N$

lvl 0



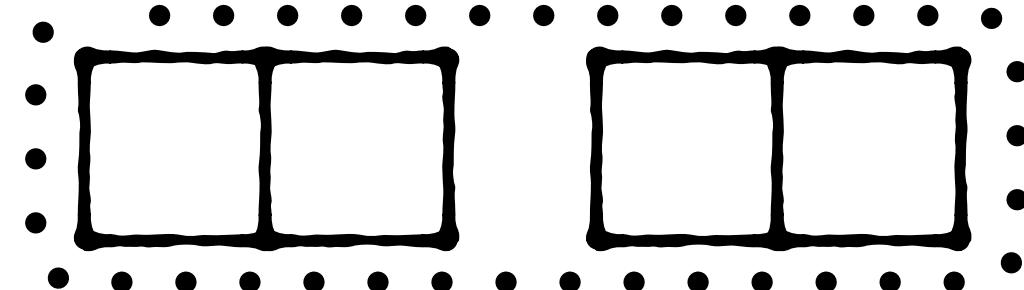
data blocks? $O(\sqrt{N})$

lvl 1

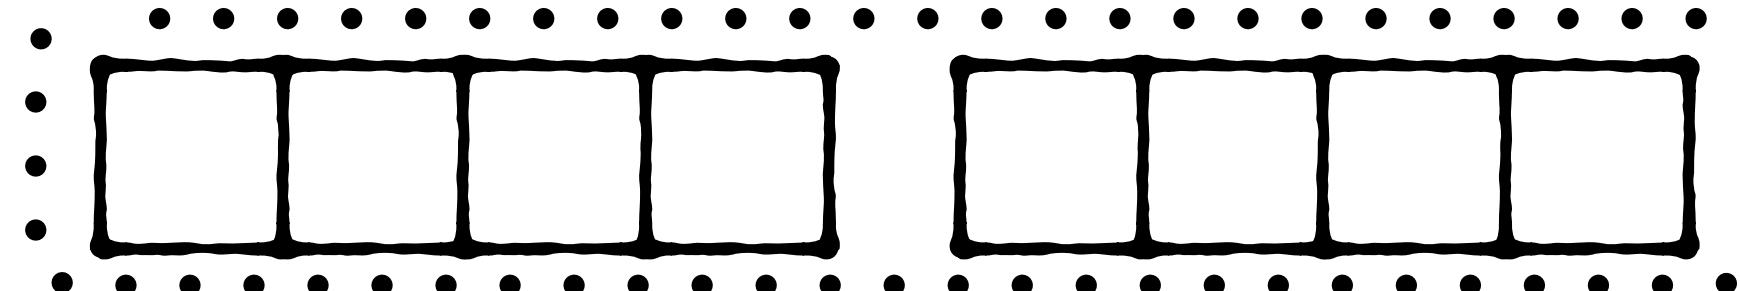


slots in largest block?

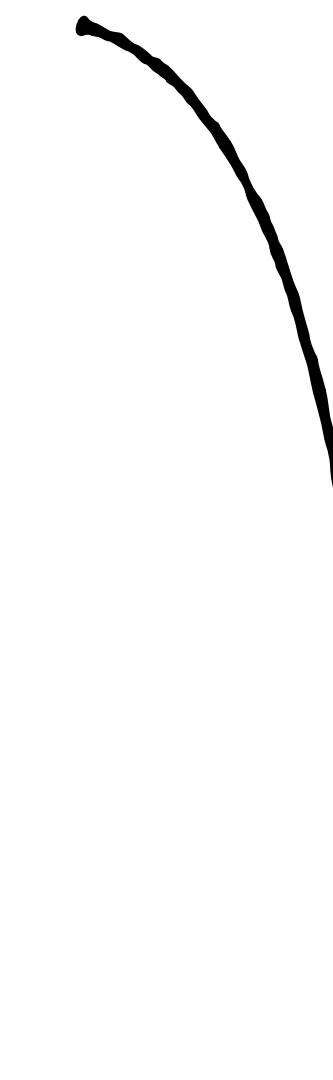
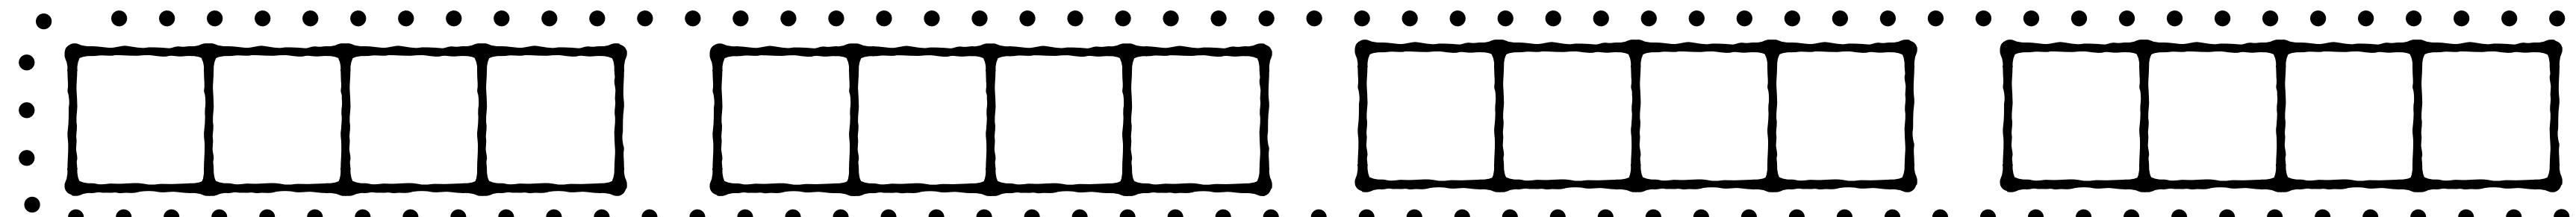
lvl 2



lvl 3



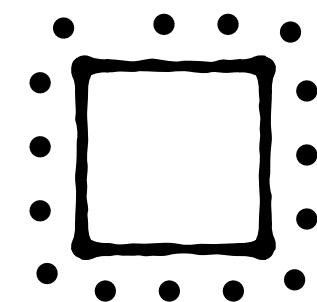
lvl 4



lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

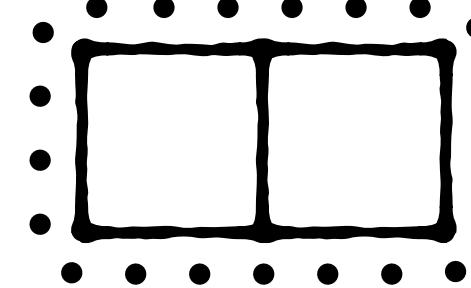
levels: $\log_2 N$

lvl 0



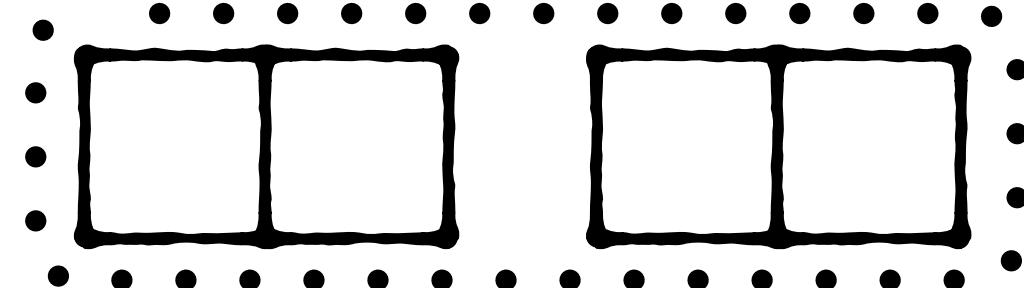
data blocks? $O(\sqrt{N})$

lvl 1

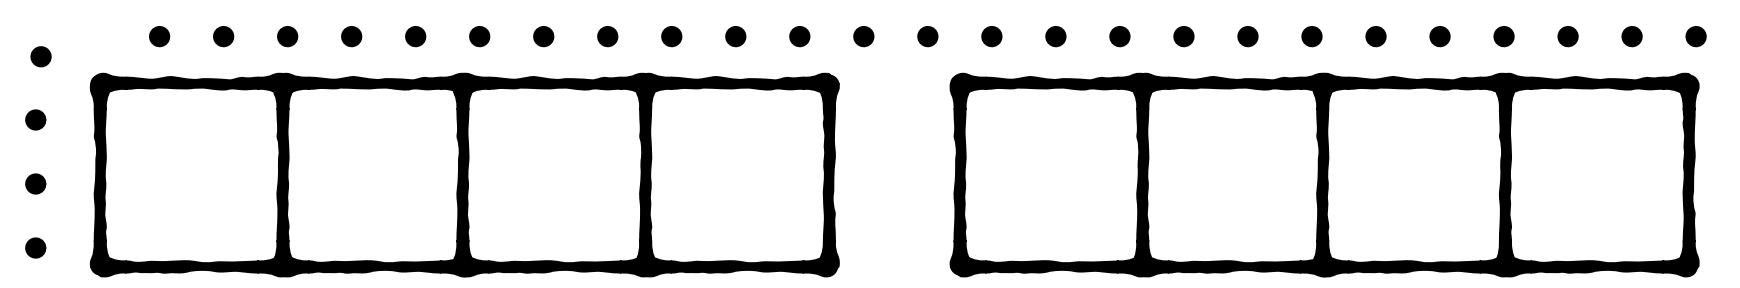


slots in largest block? $2^{\lceil K/2 \rceil}$

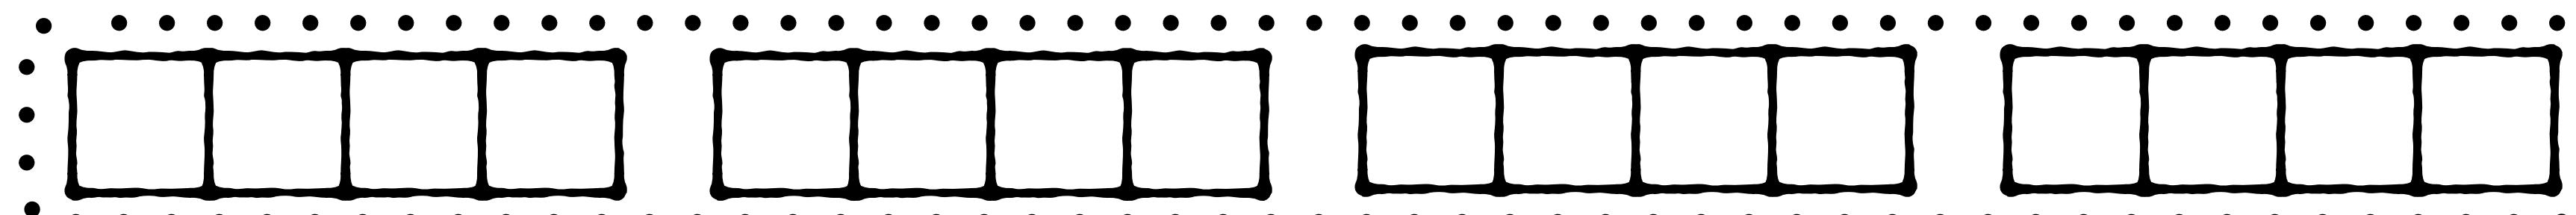
lvl 2



lvl 3



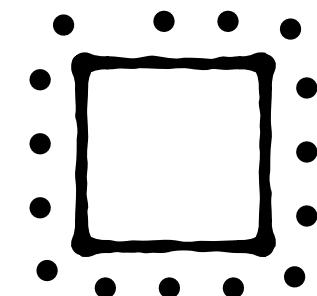
lvl 4



lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

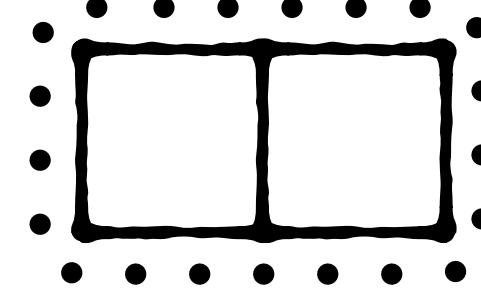
levels: $\log_2 N$

lvl 0



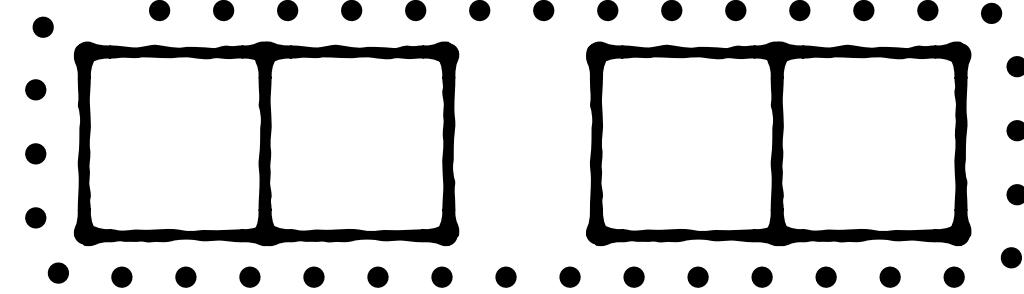
data blocks? $O(\sqrt{N})$

lvl 1

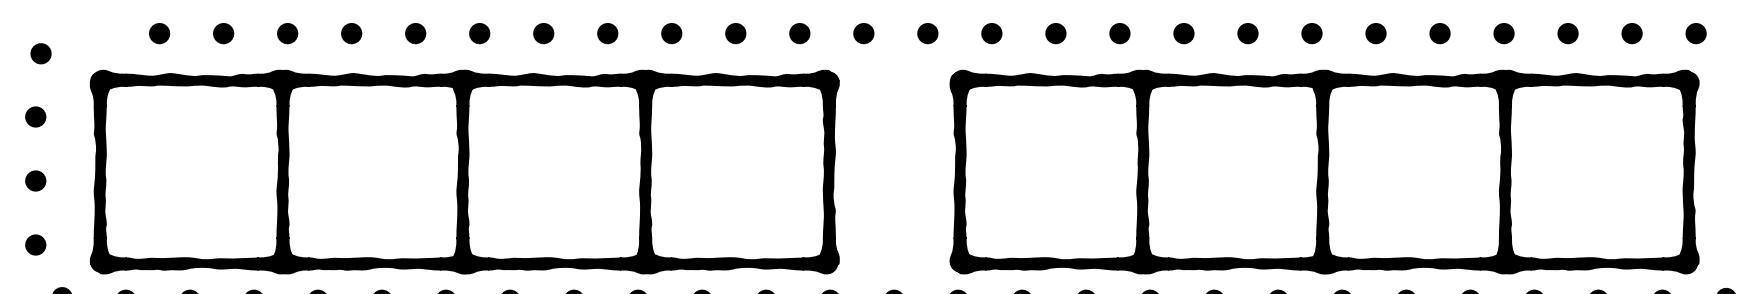


slots in largest block? $O(\sqrt{N})$

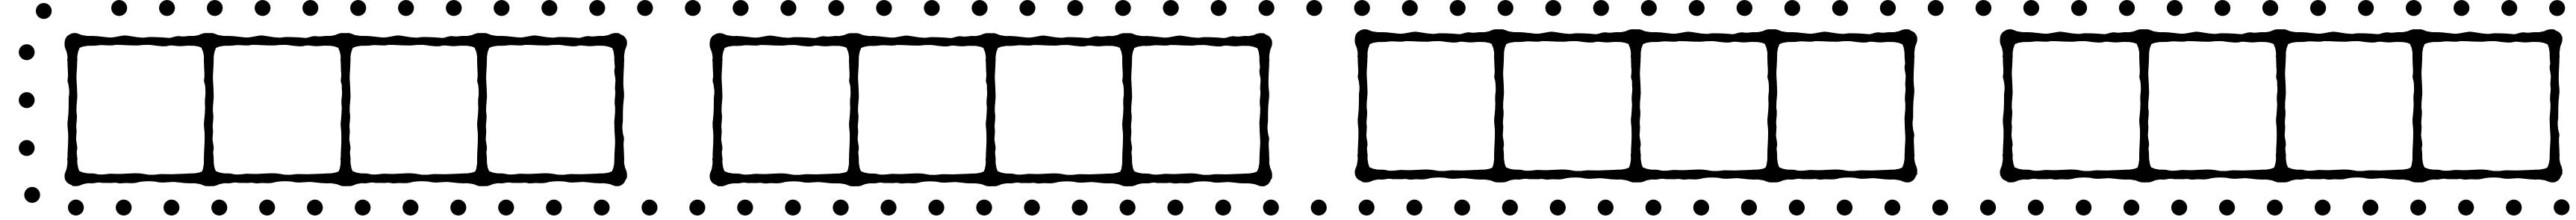
lvl 2



lvl 3



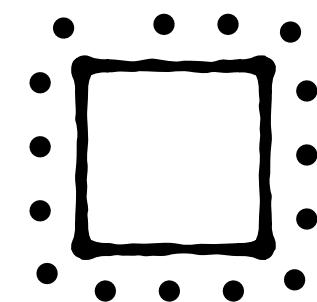
lvl 4



lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

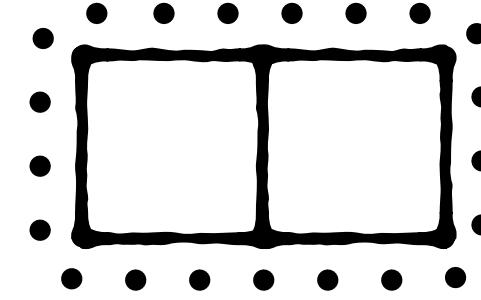
levels: $\log_2 N$

lvl 0

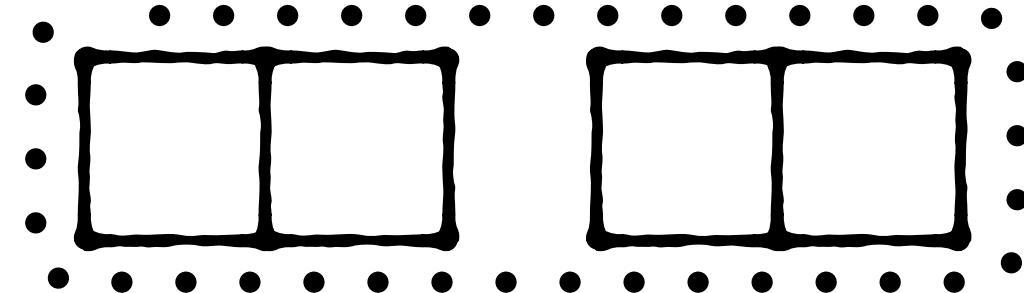


data blocks? $O(\sqrt{N})$

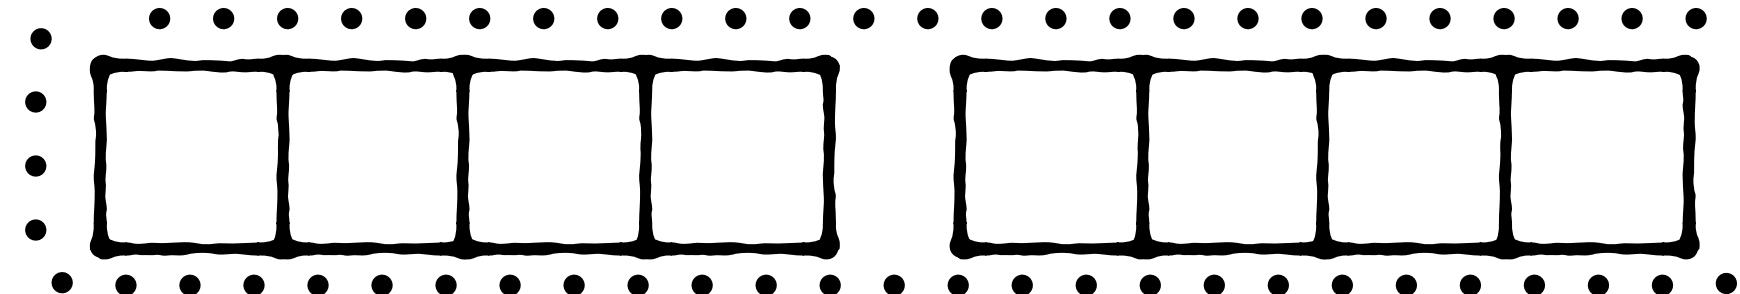
lvl 1



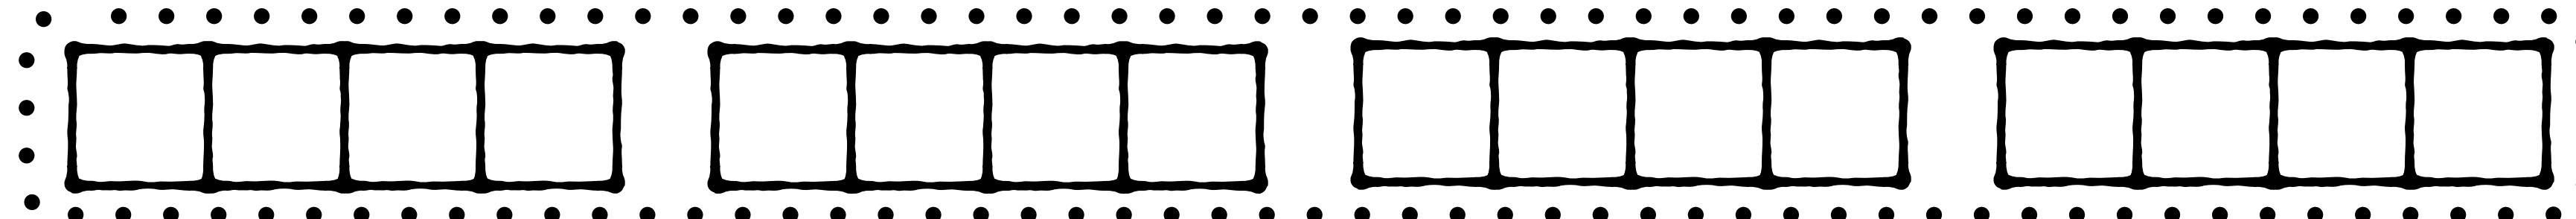
lvl 2



lvl 3



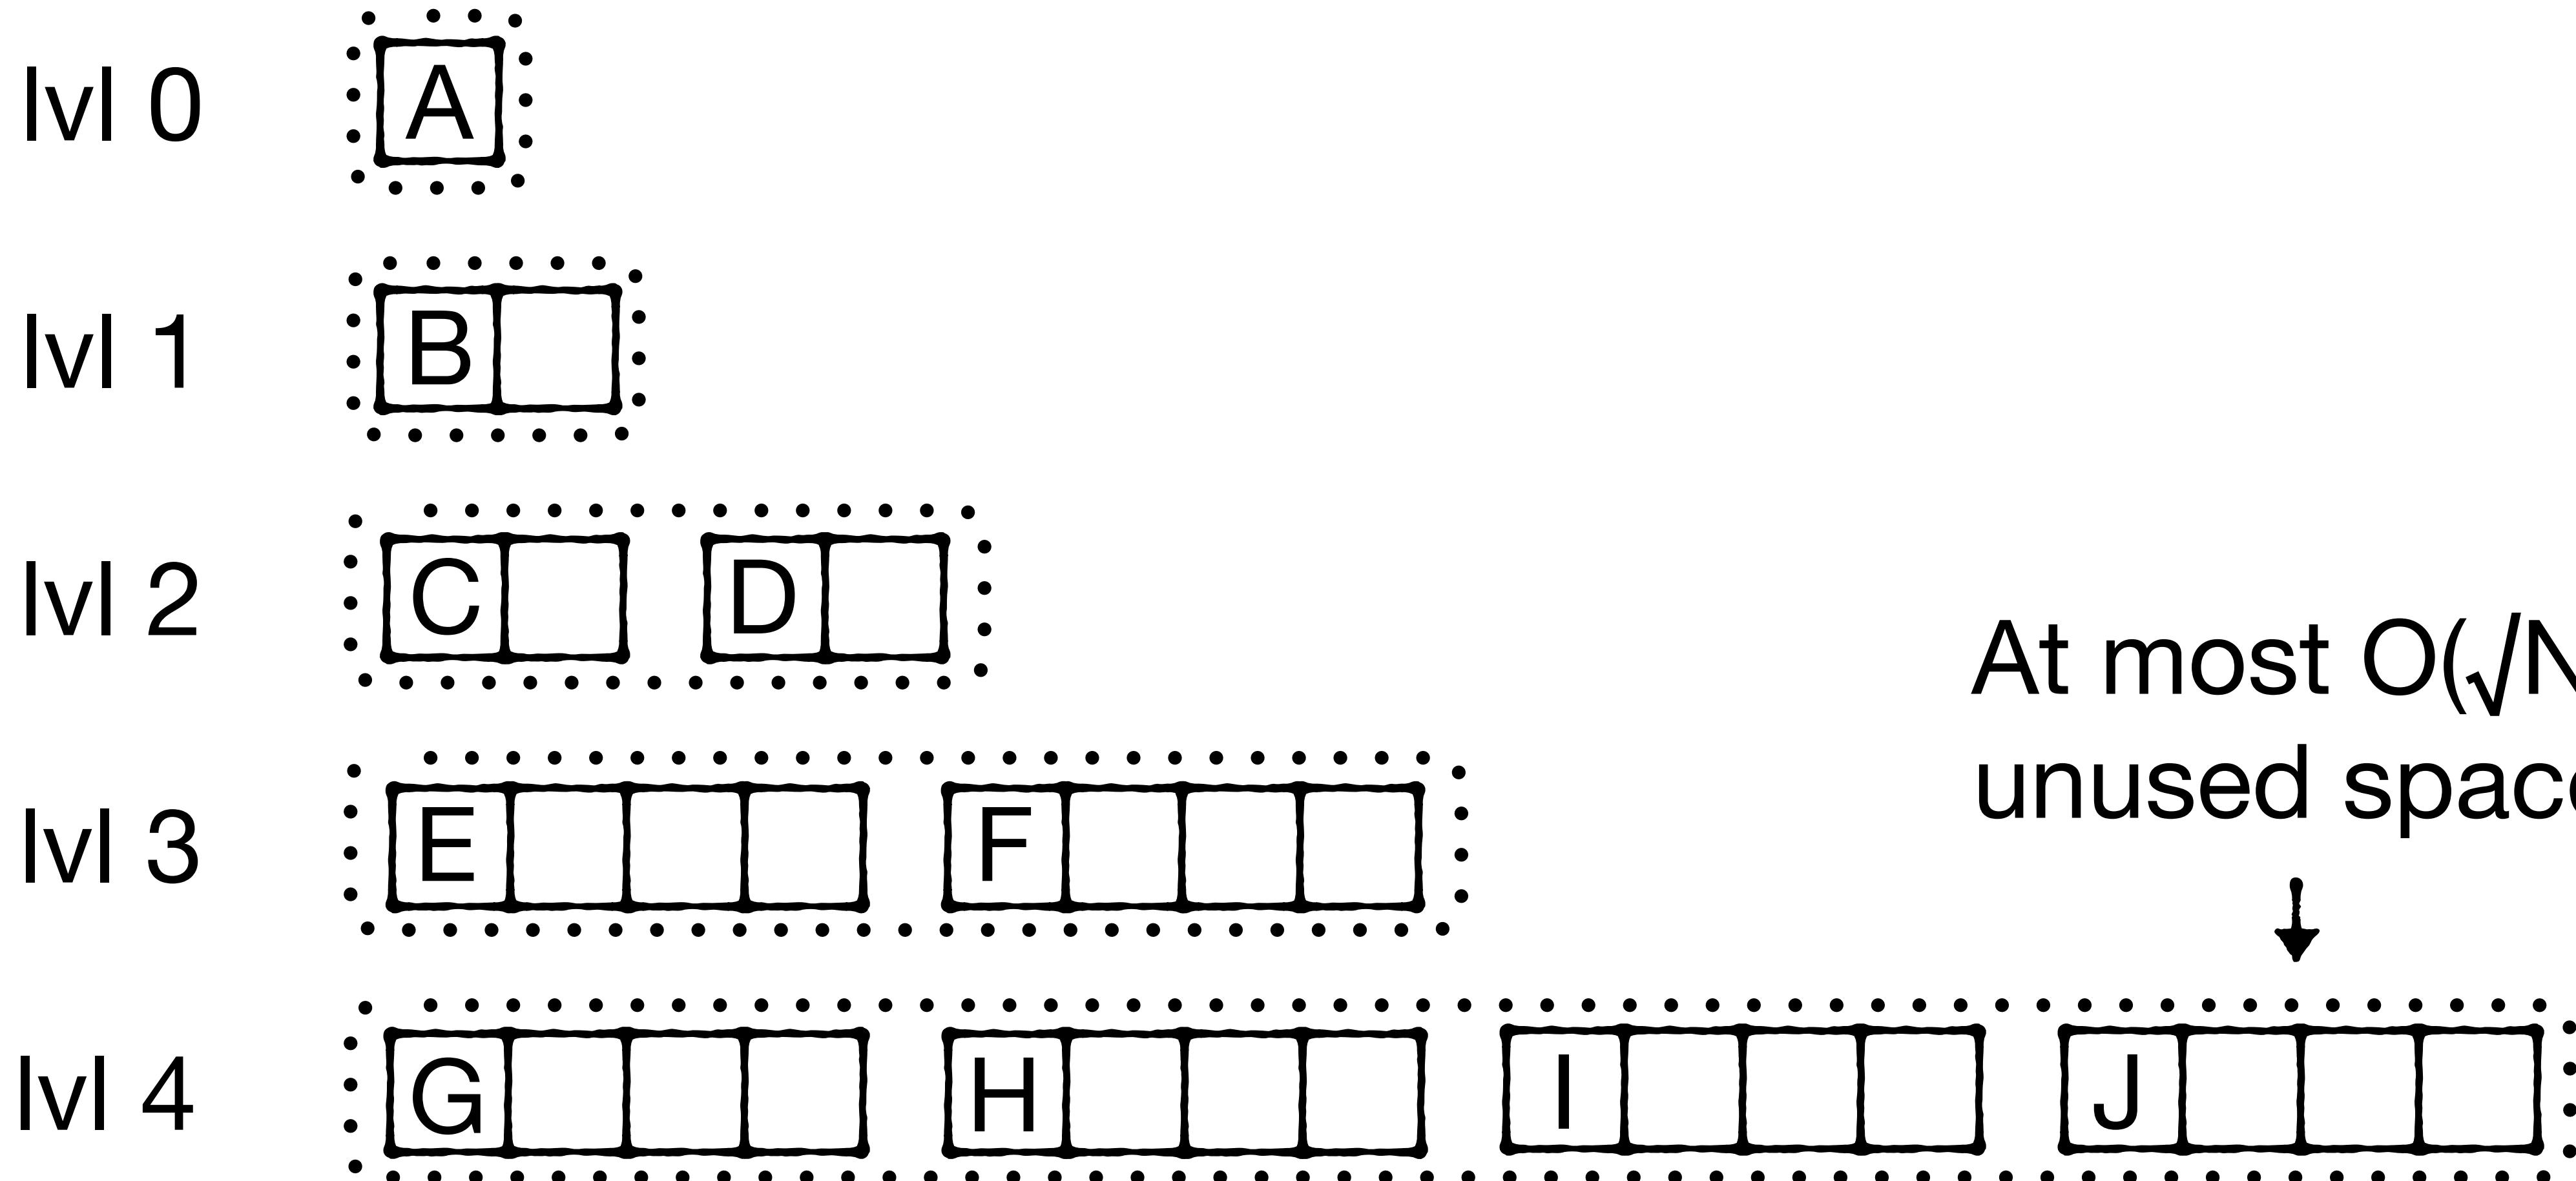
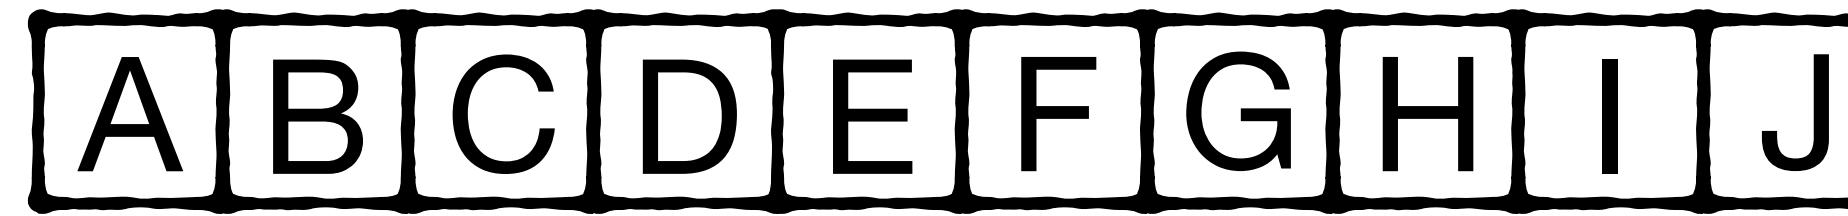
lvl 4



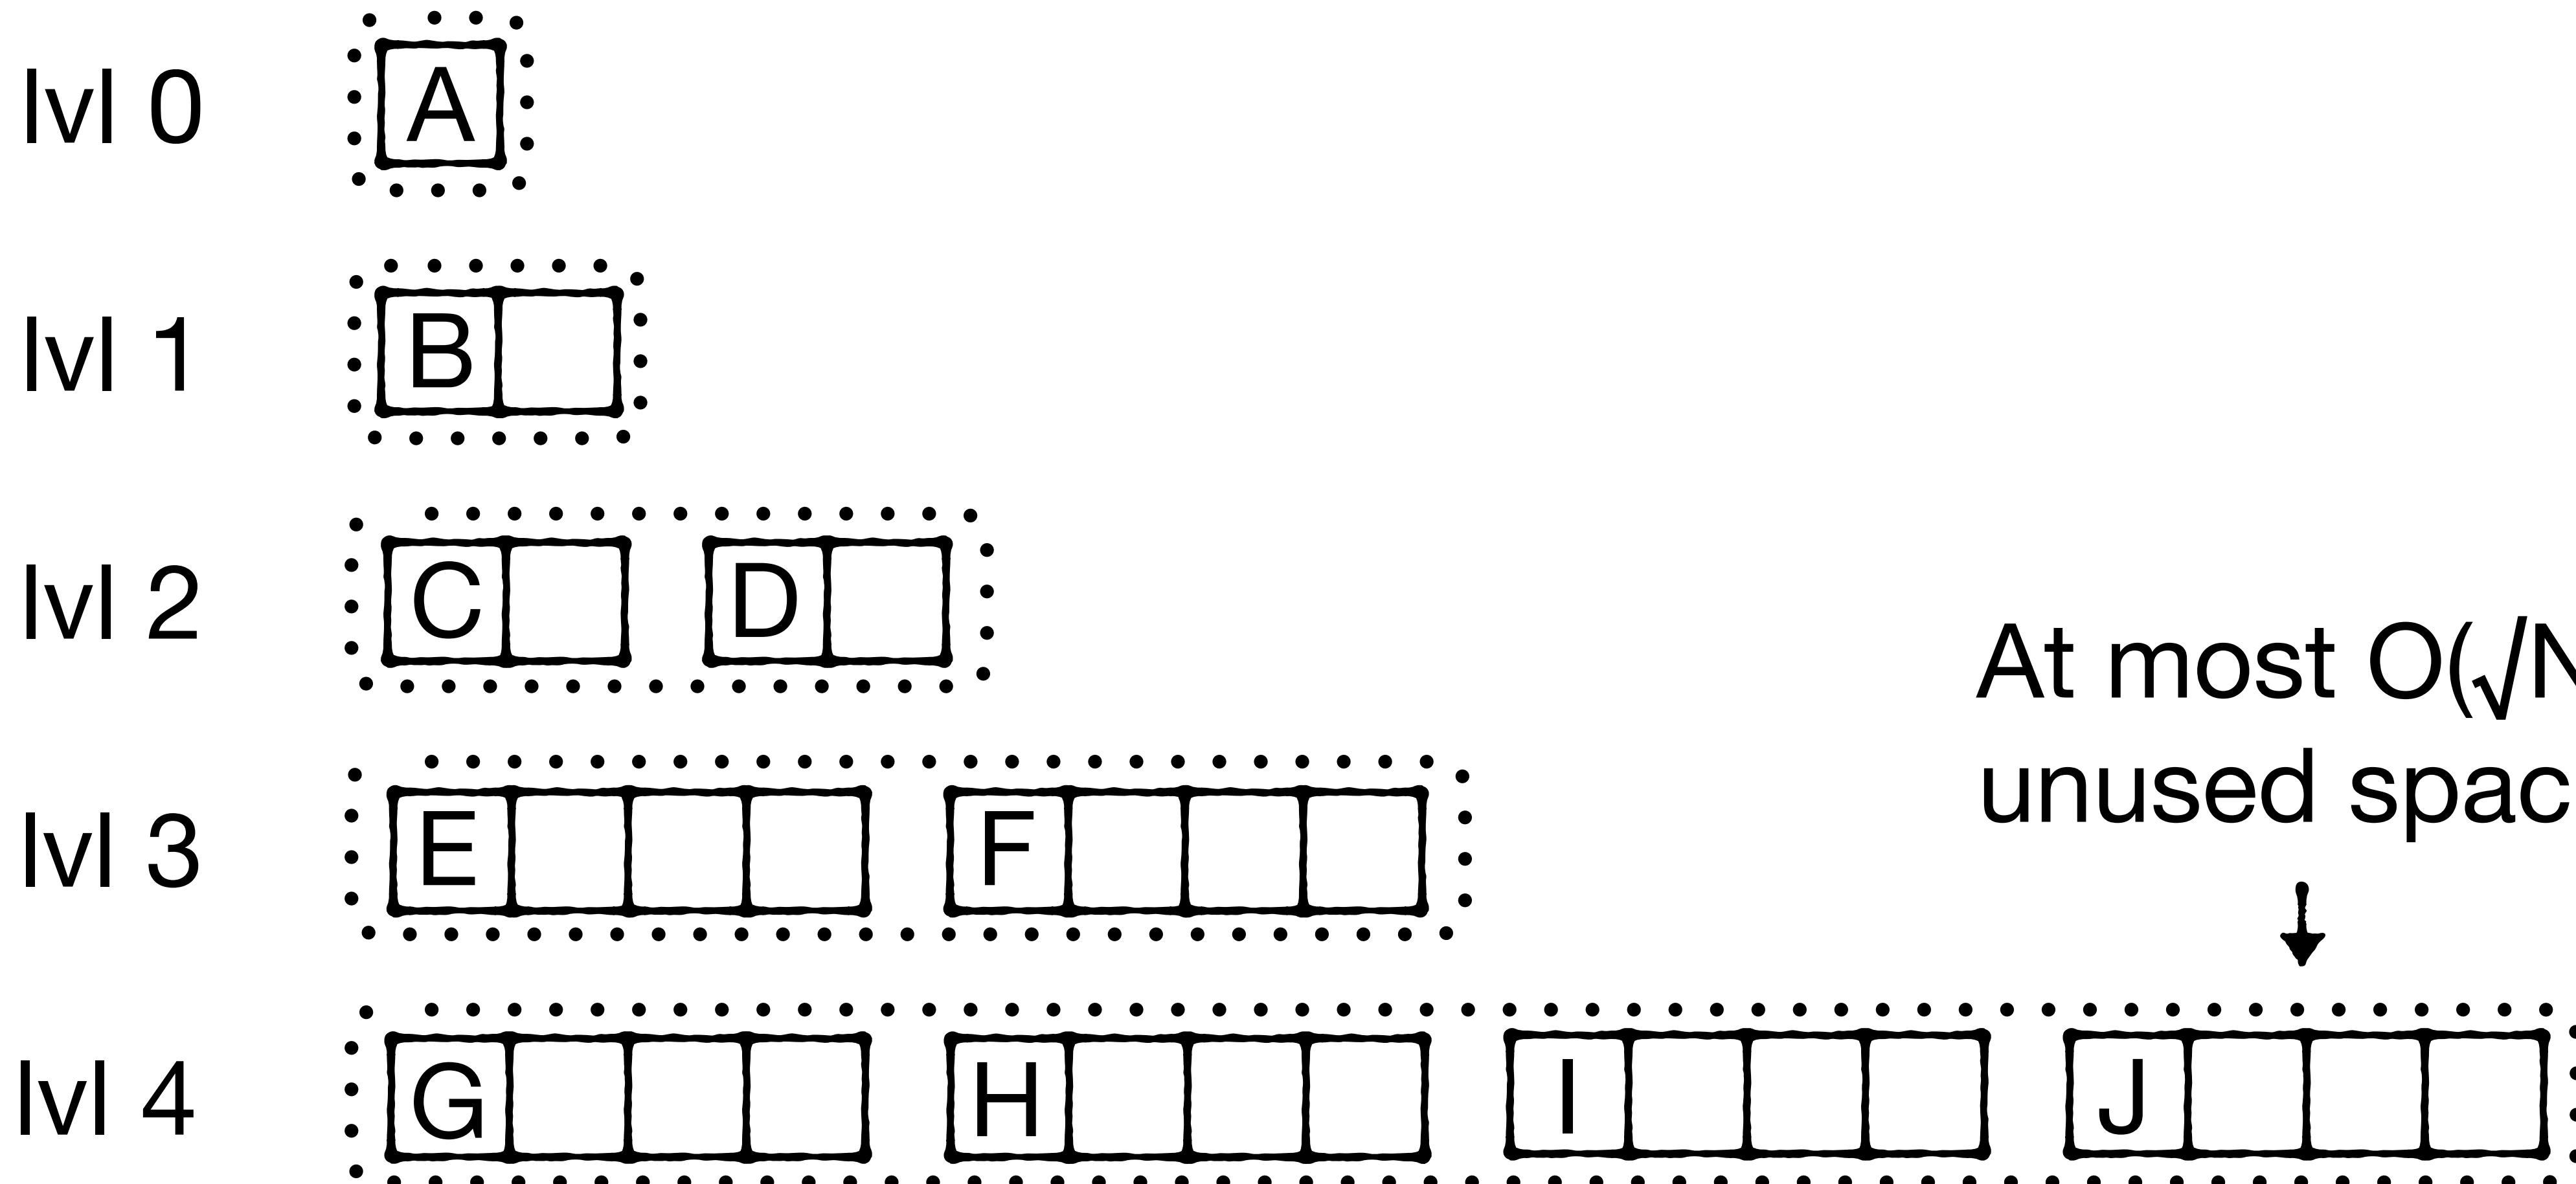
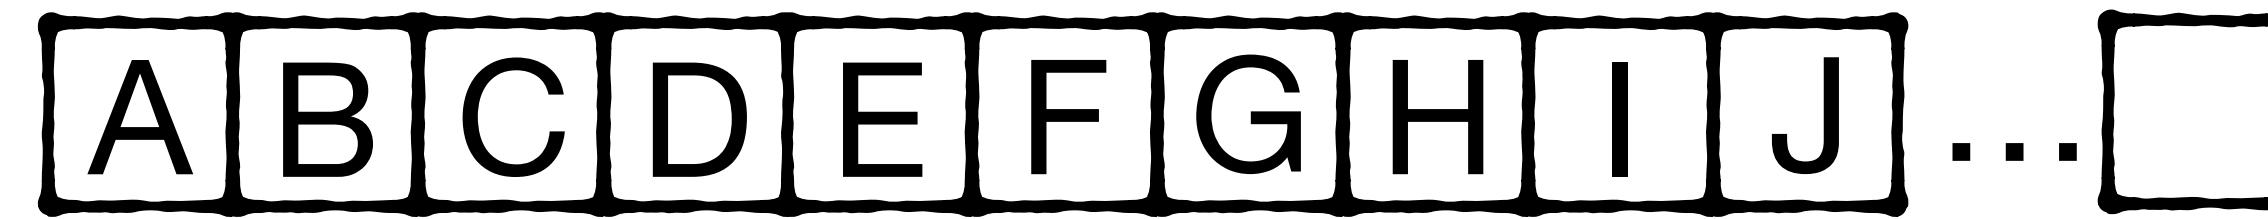
At most $O(\sqrt{N})$
unused space



Directory with $O(\sqrt{N})$ pointers

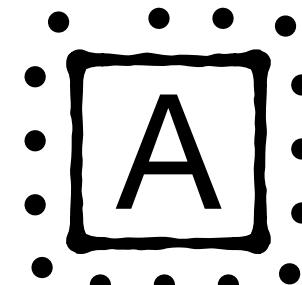


At most half $O(\sqrt{N})$ unused space

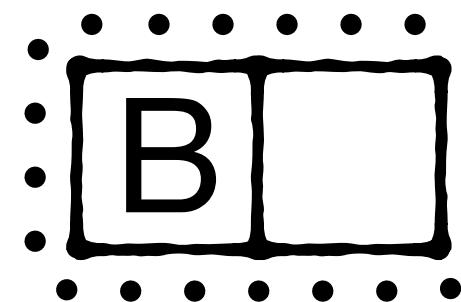


A B C D E F G H I J ...

lvl 0

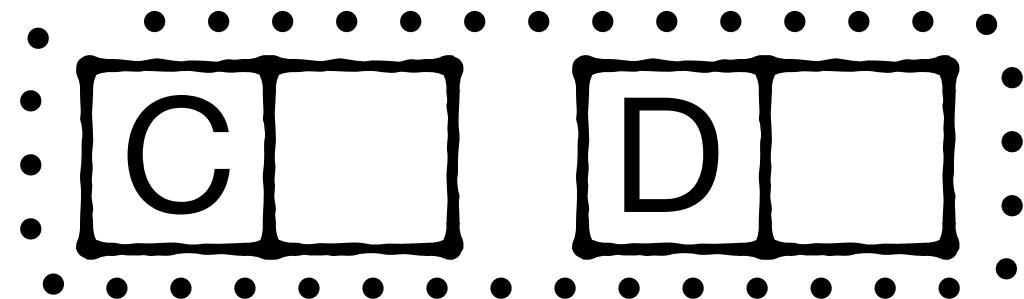


lvl 1

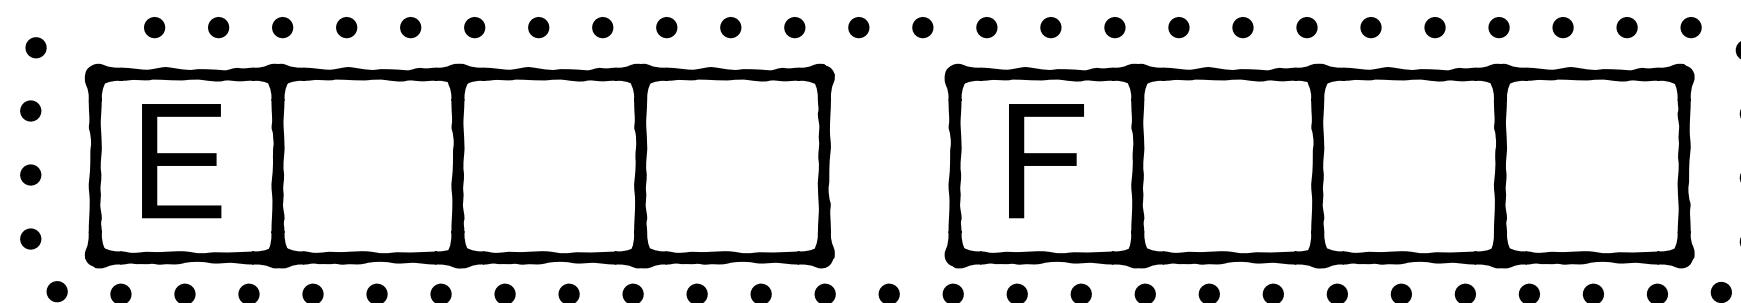


Max extra space: $O(\sqrt{N}) + O(\sqrt{N}) = O(\sqrt{N})$

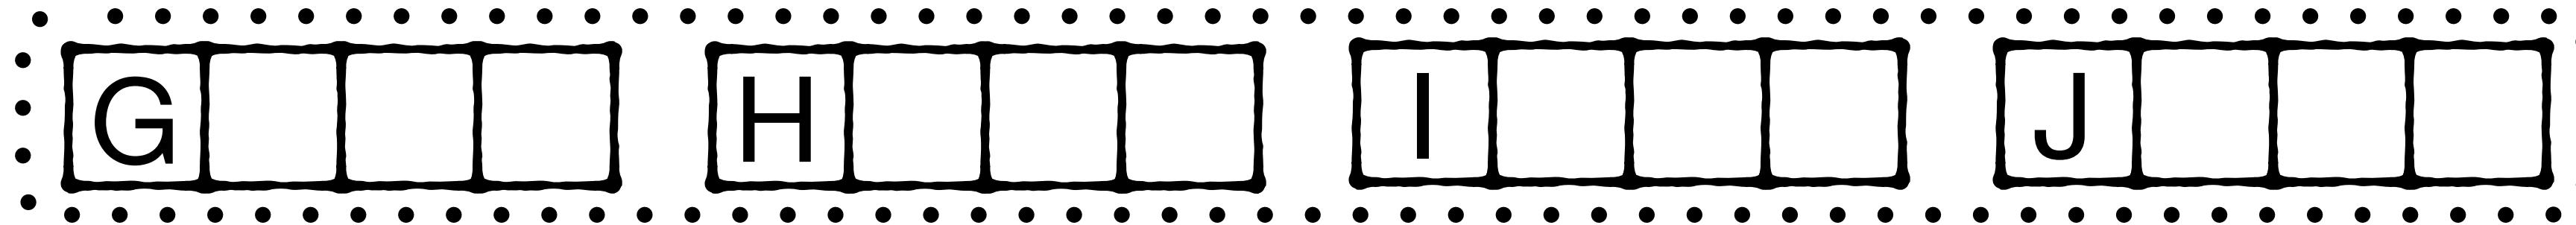
lvl 2



lvl 3



lvl 4



A B C D E F G H I J ...

lvl 0

A

lvl 1

B

lvl 2

C D

lvl 3

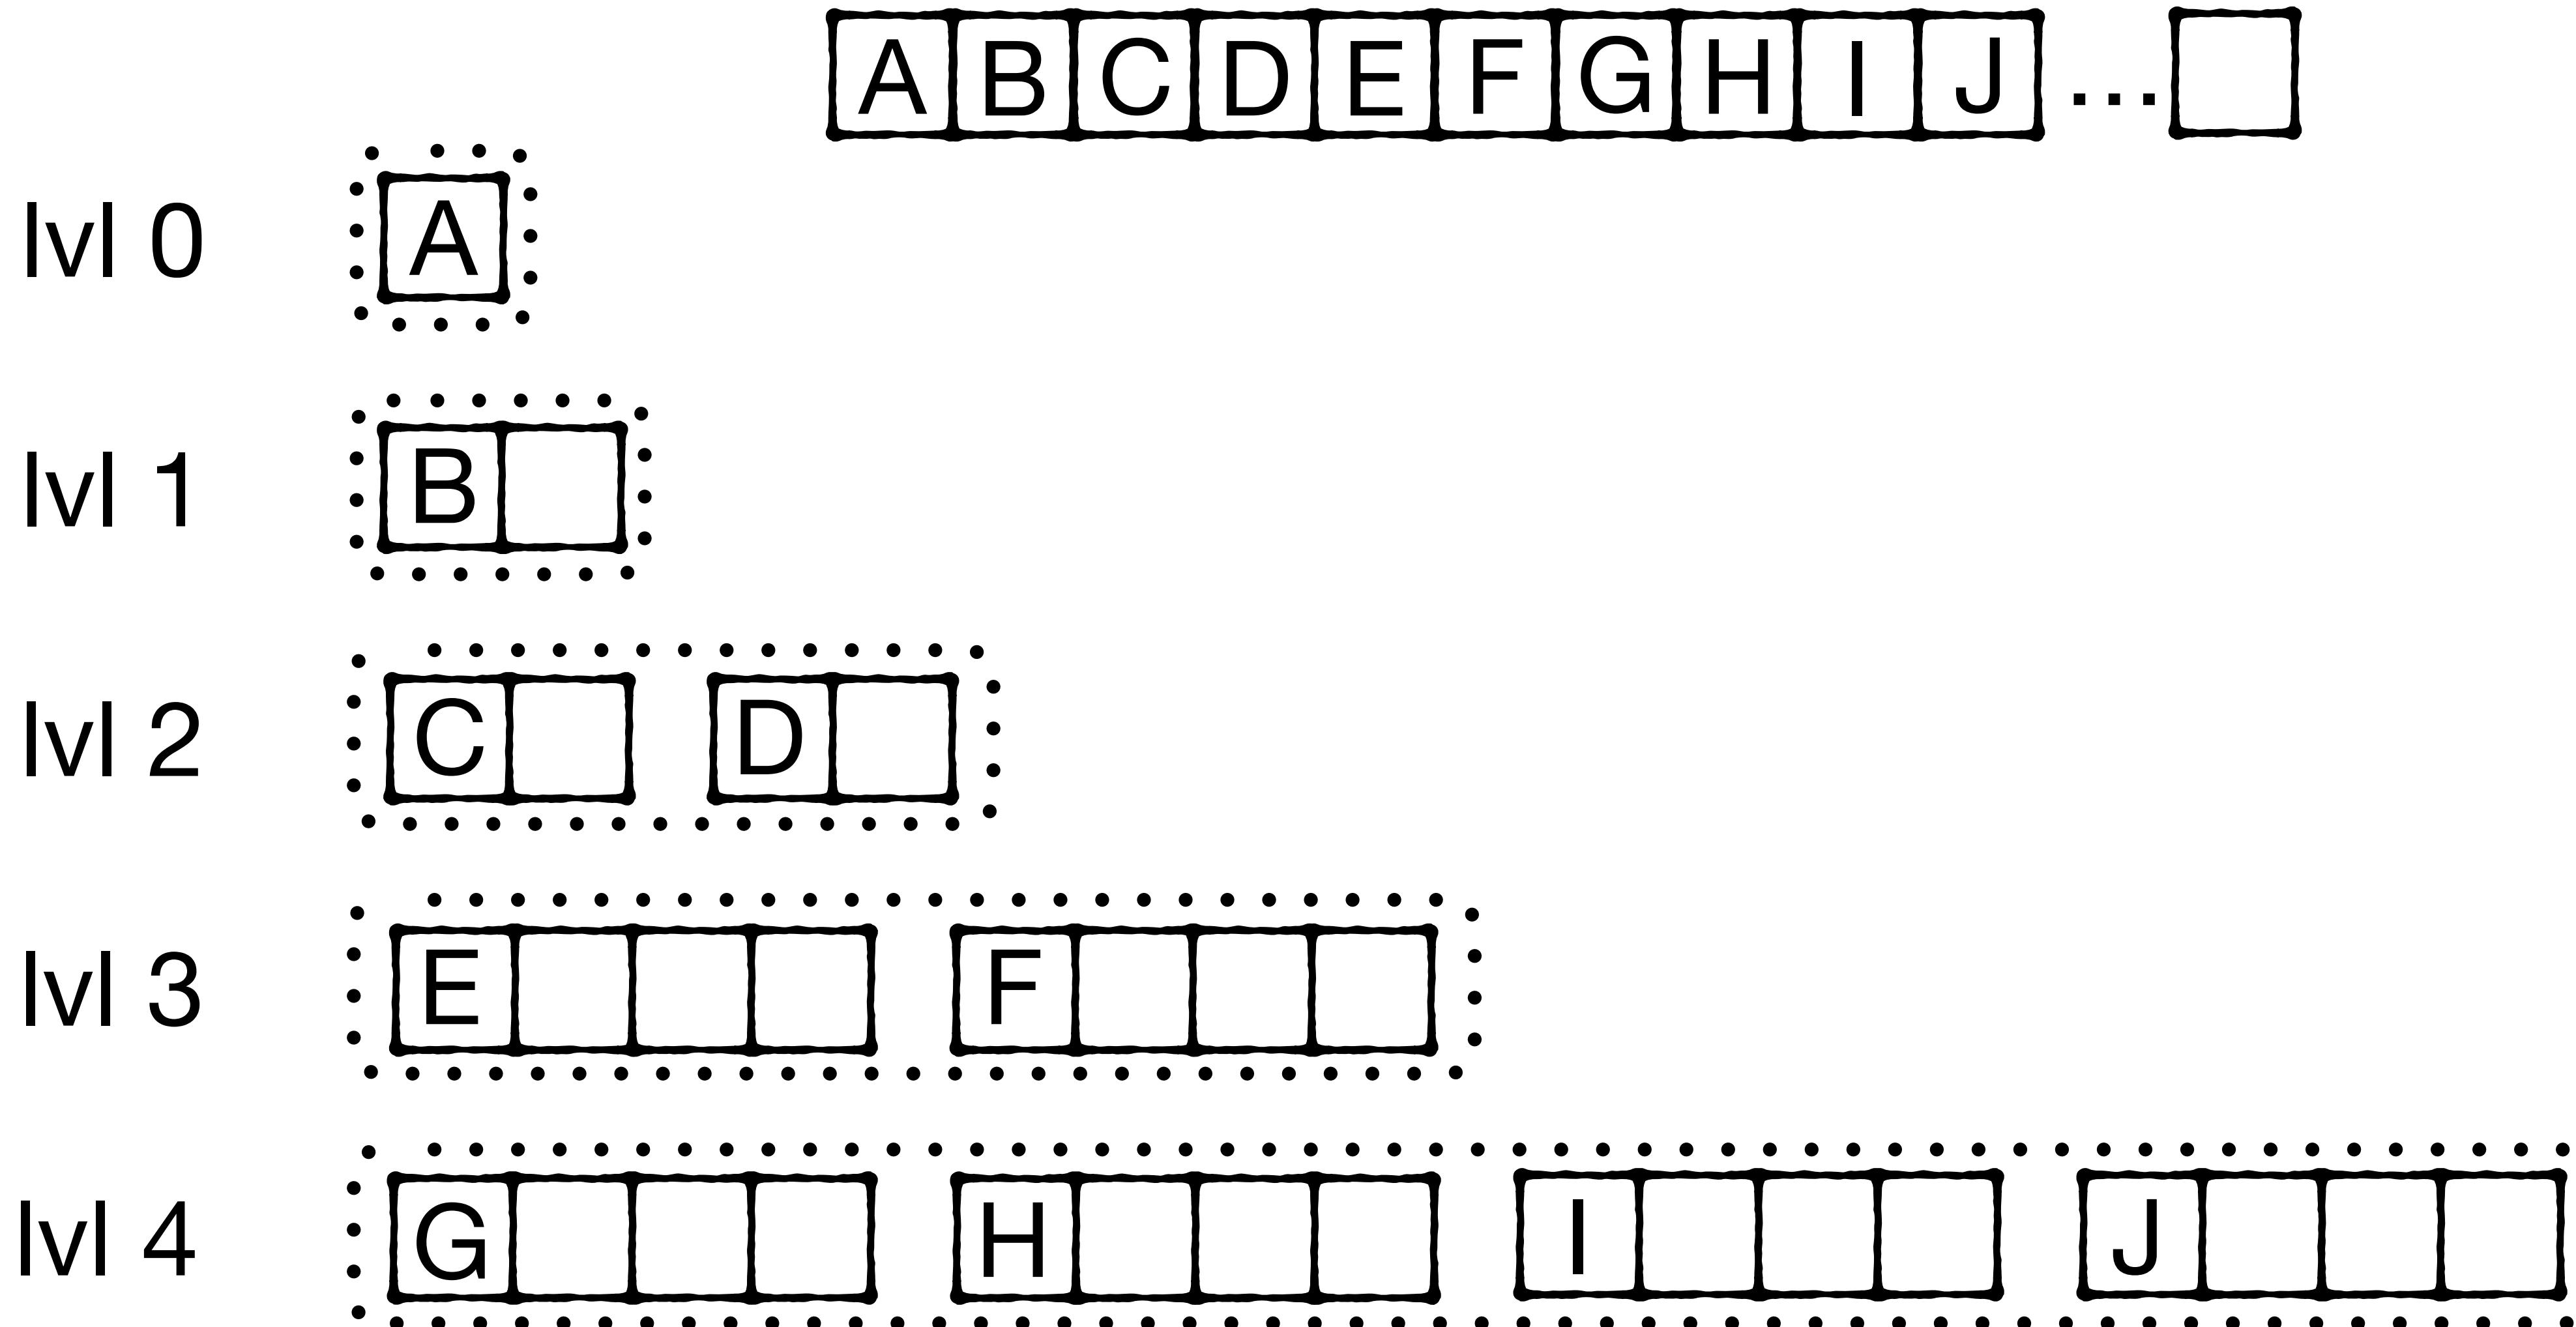
E F

lvl 4

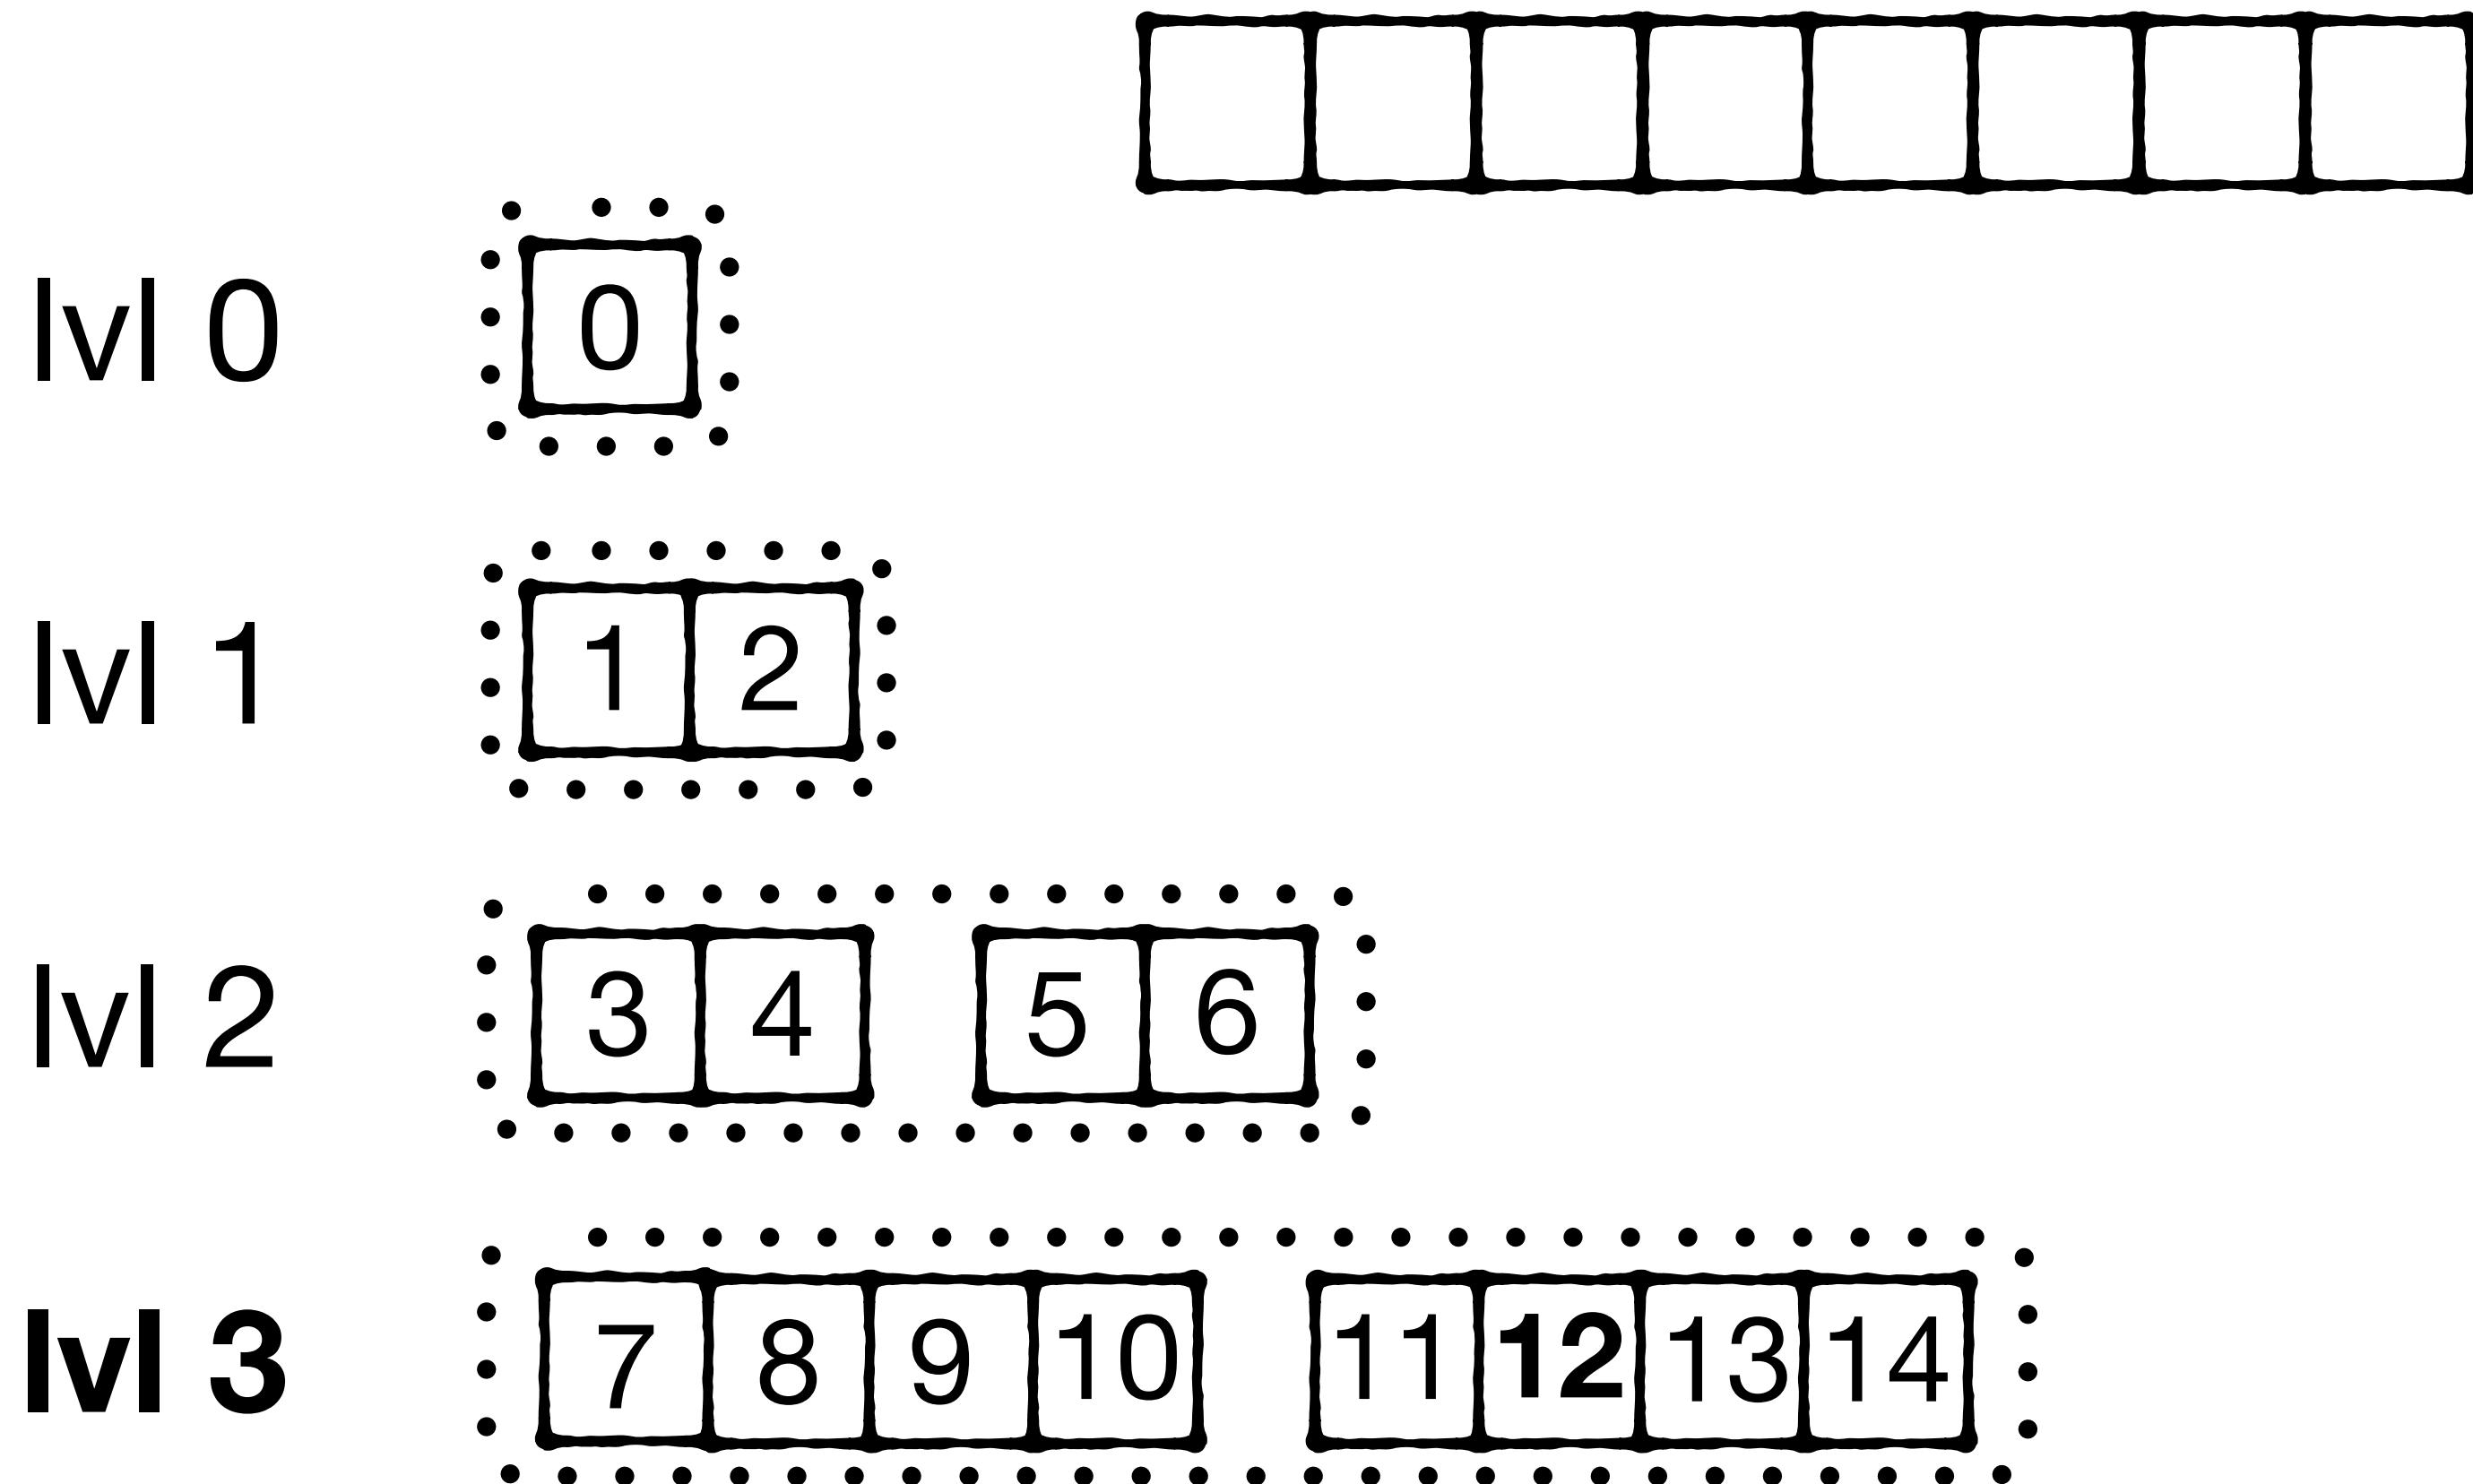
G H I J

Max space-amp: = $O(1+1/\sqrt{N})$

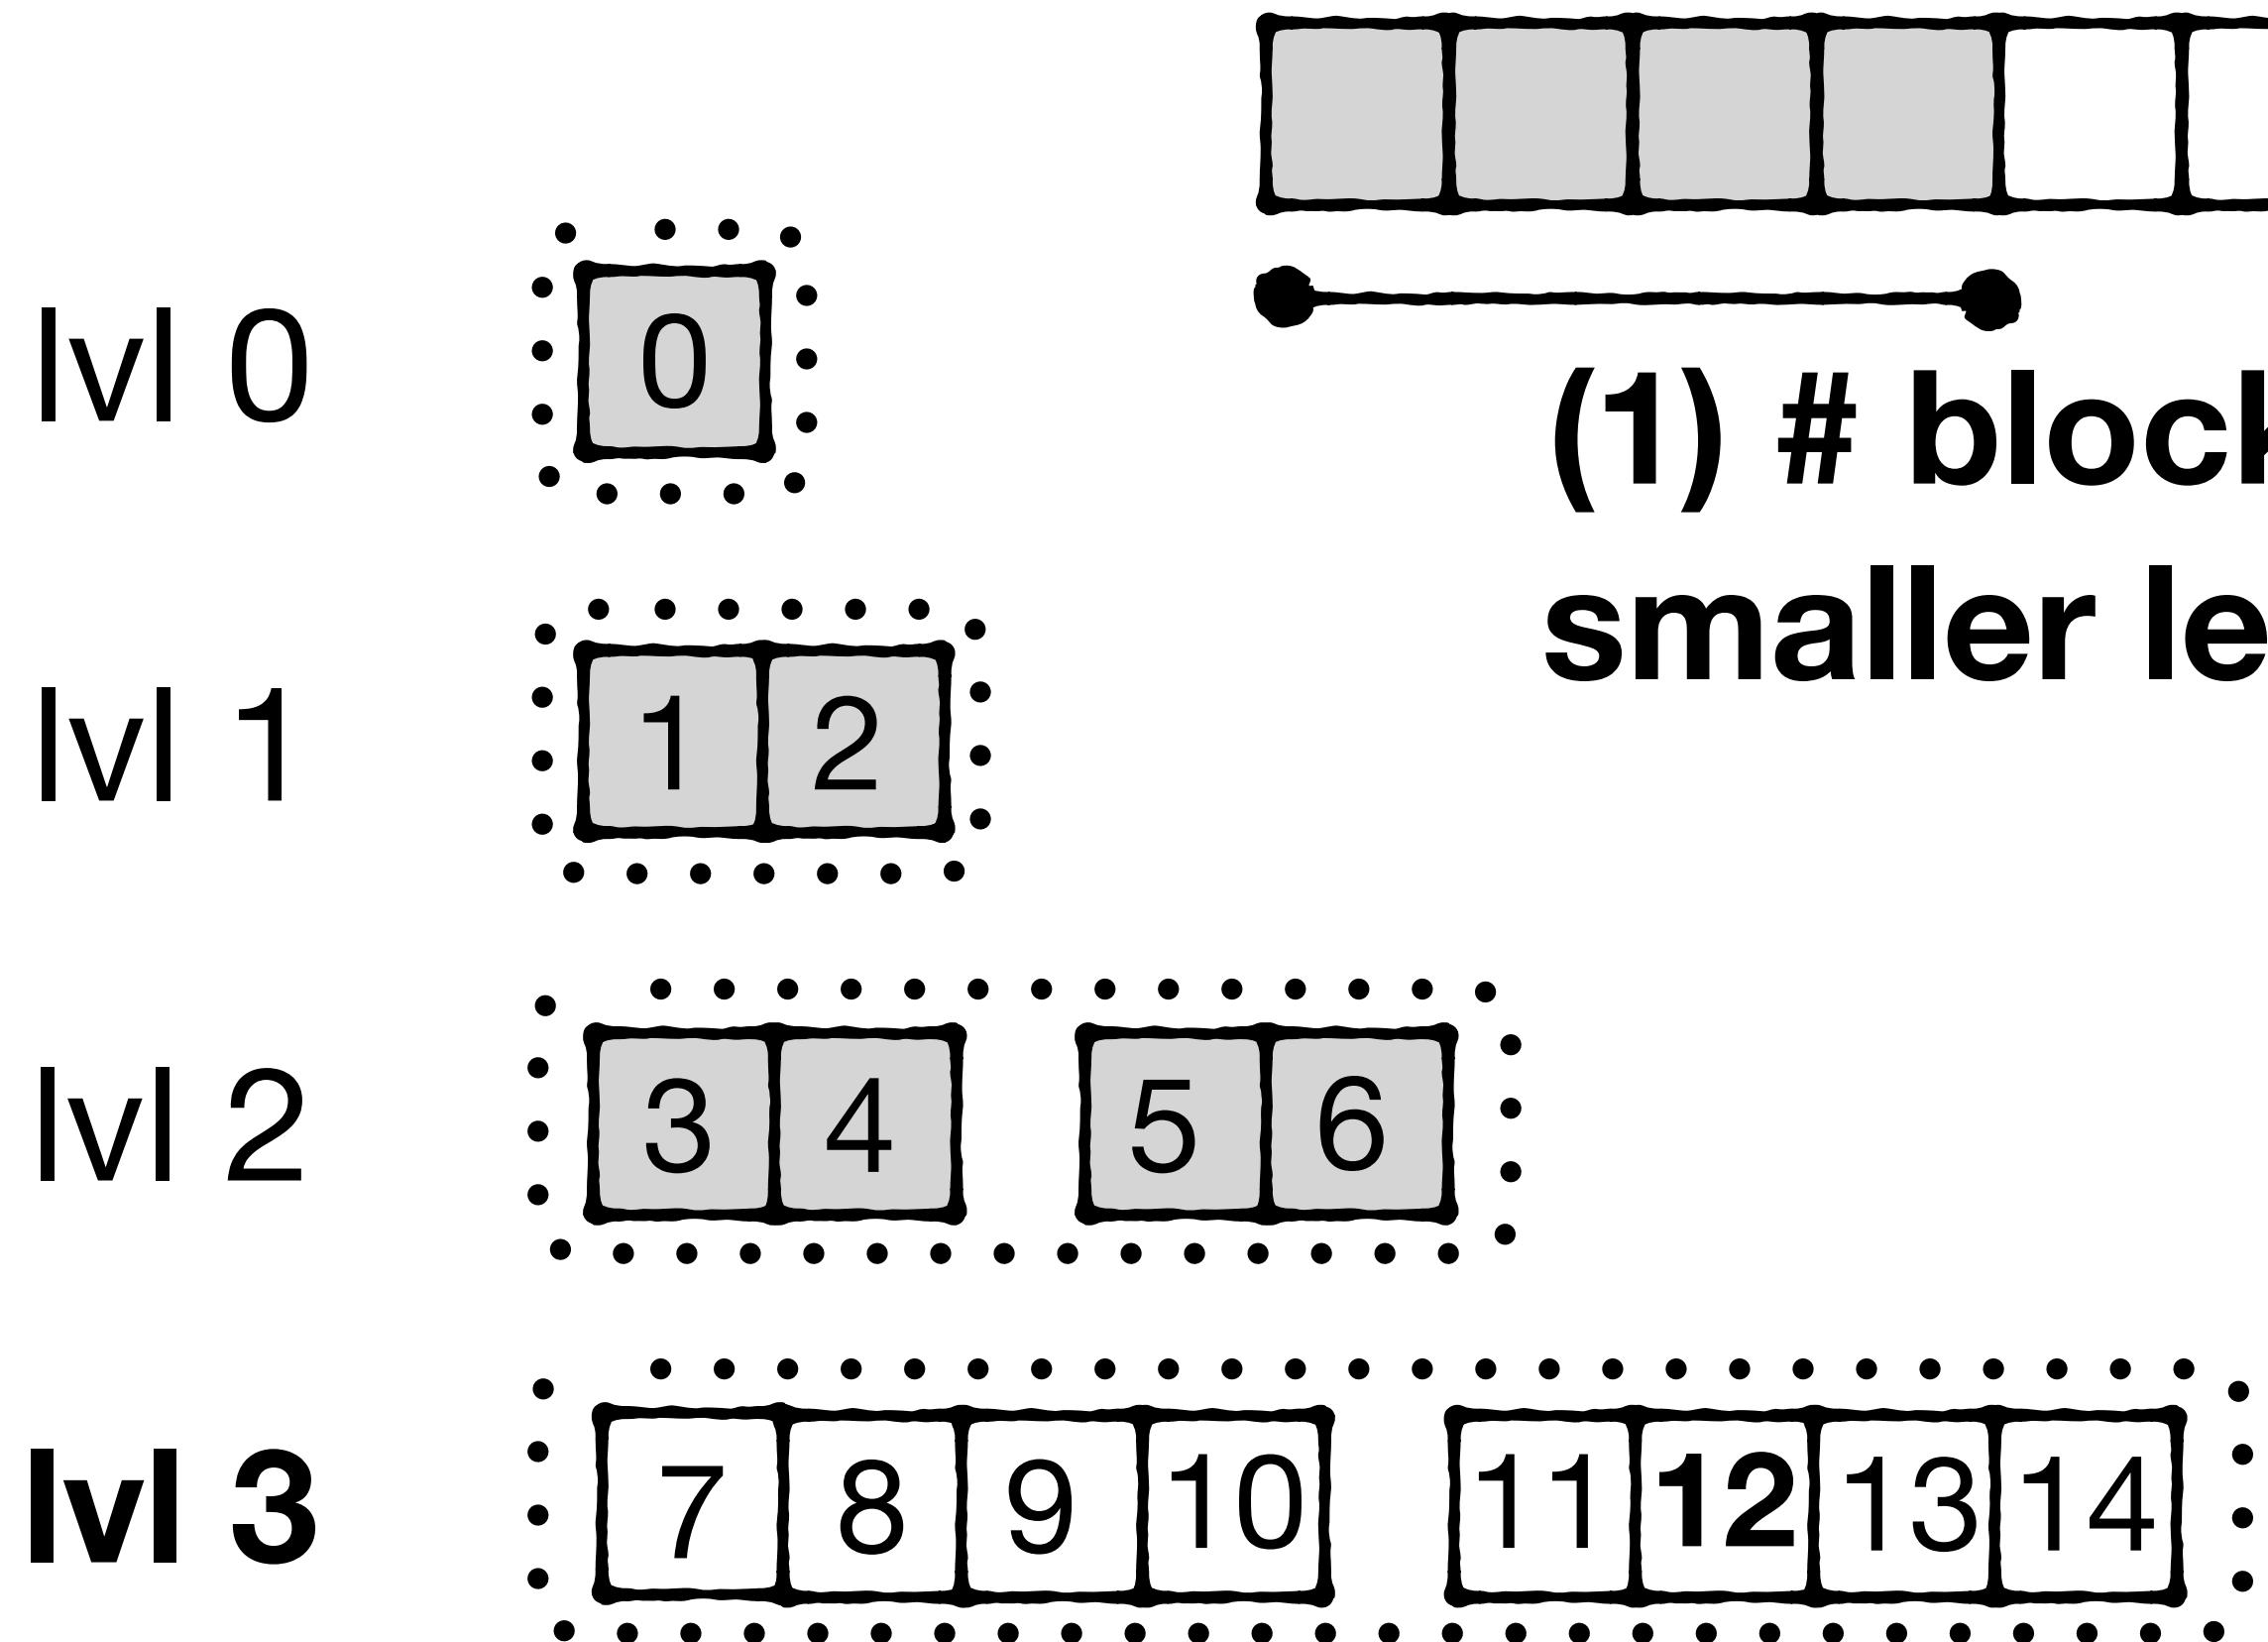
How to access slot in O(1) time?



get(12)

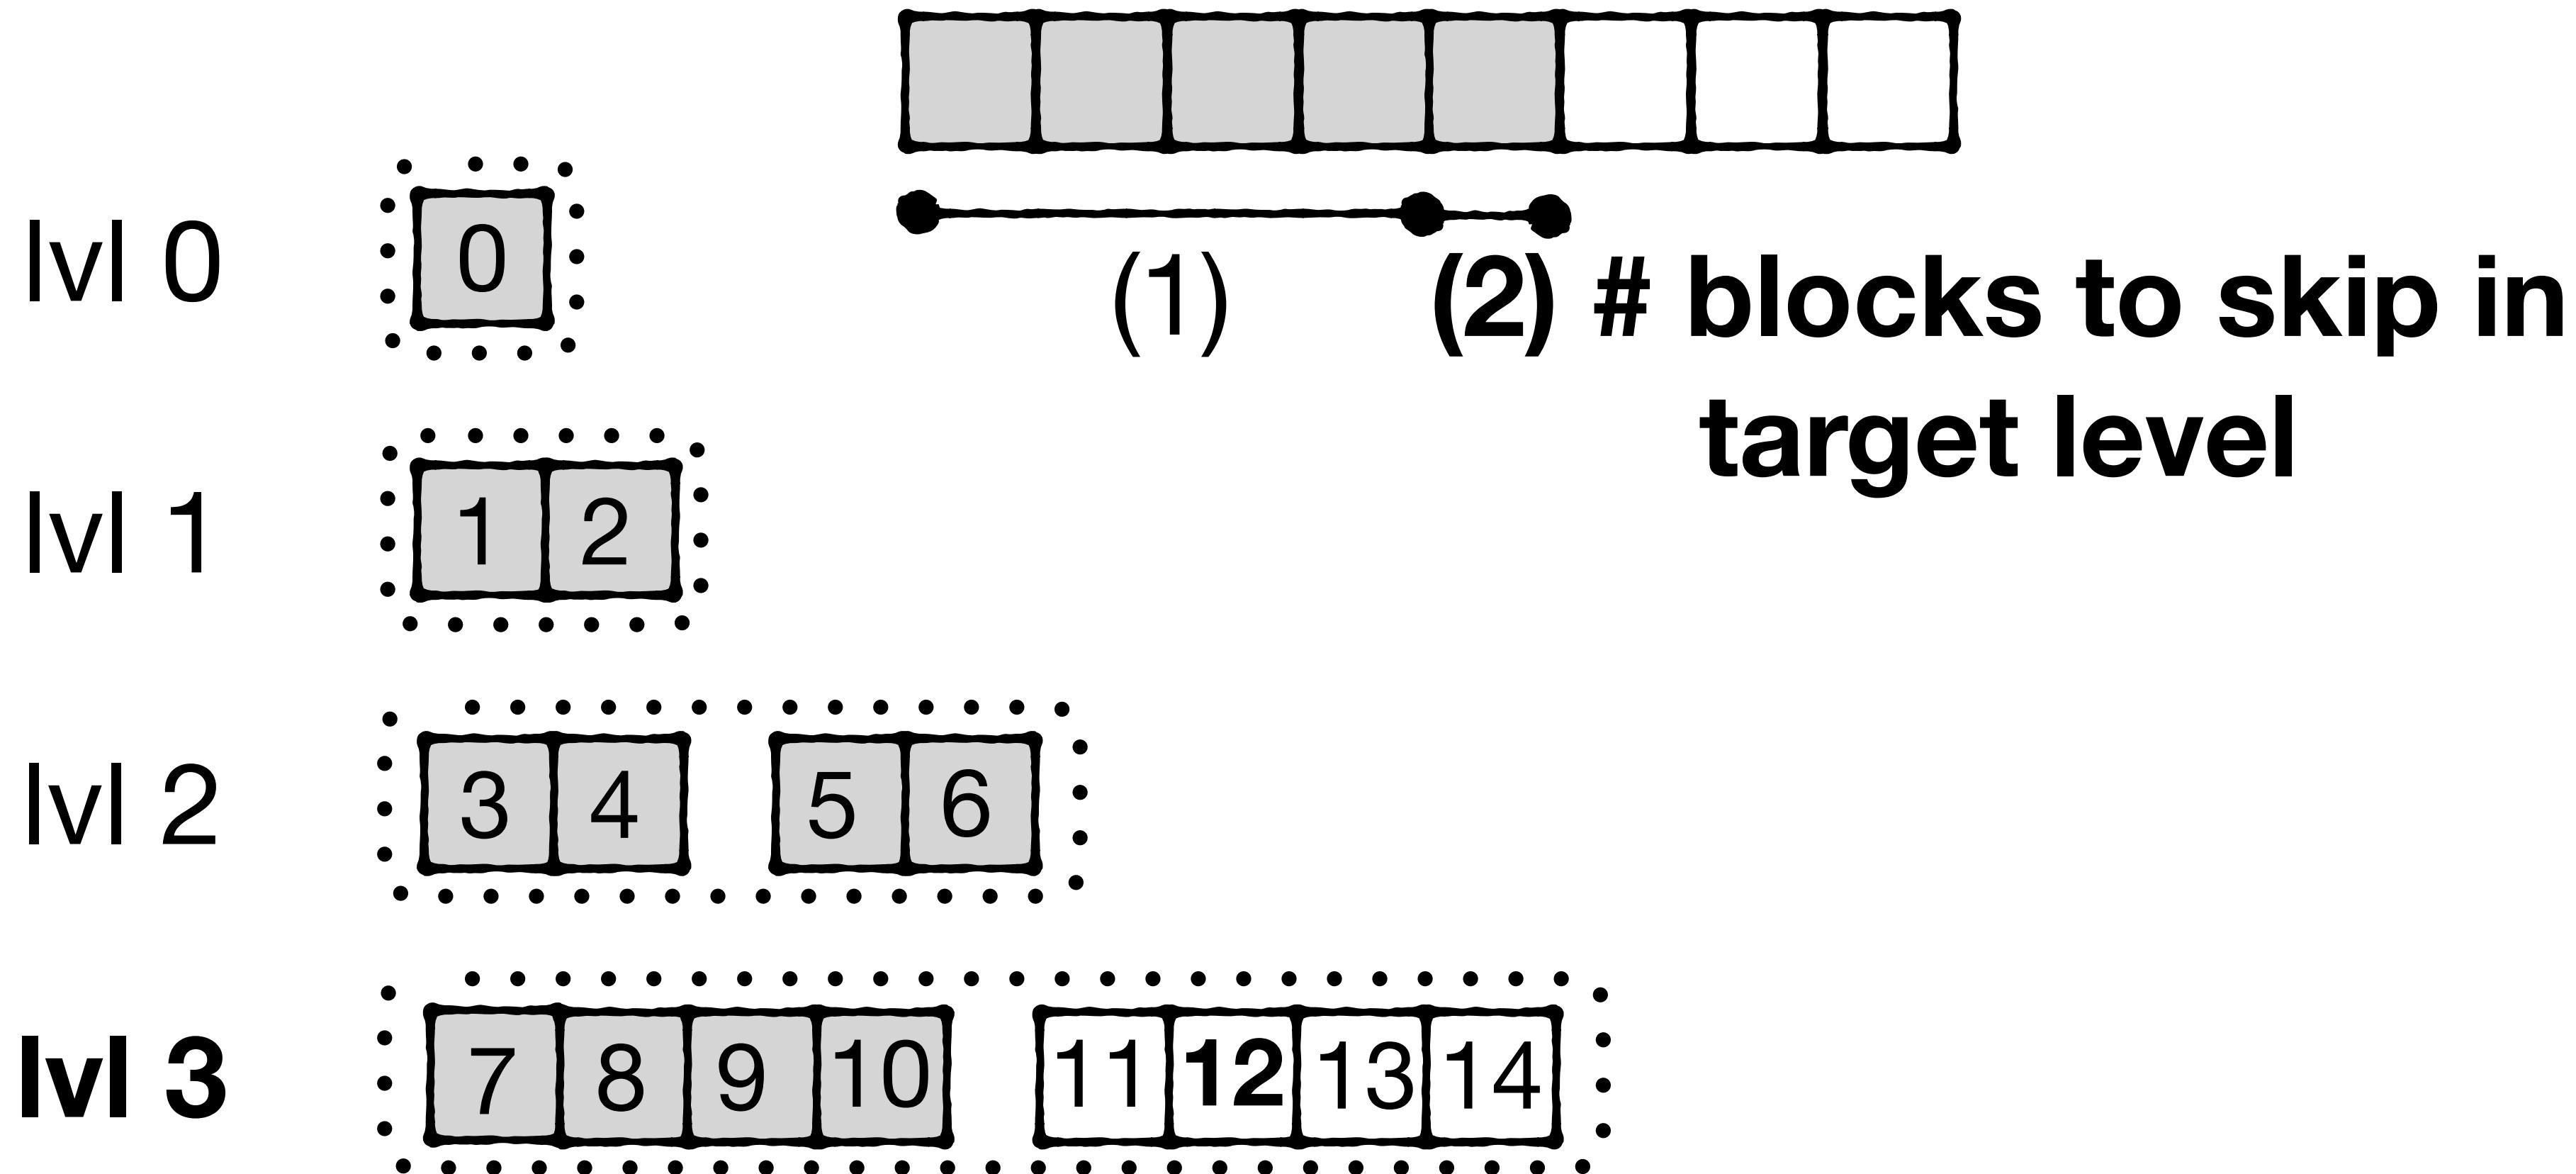


get(12)

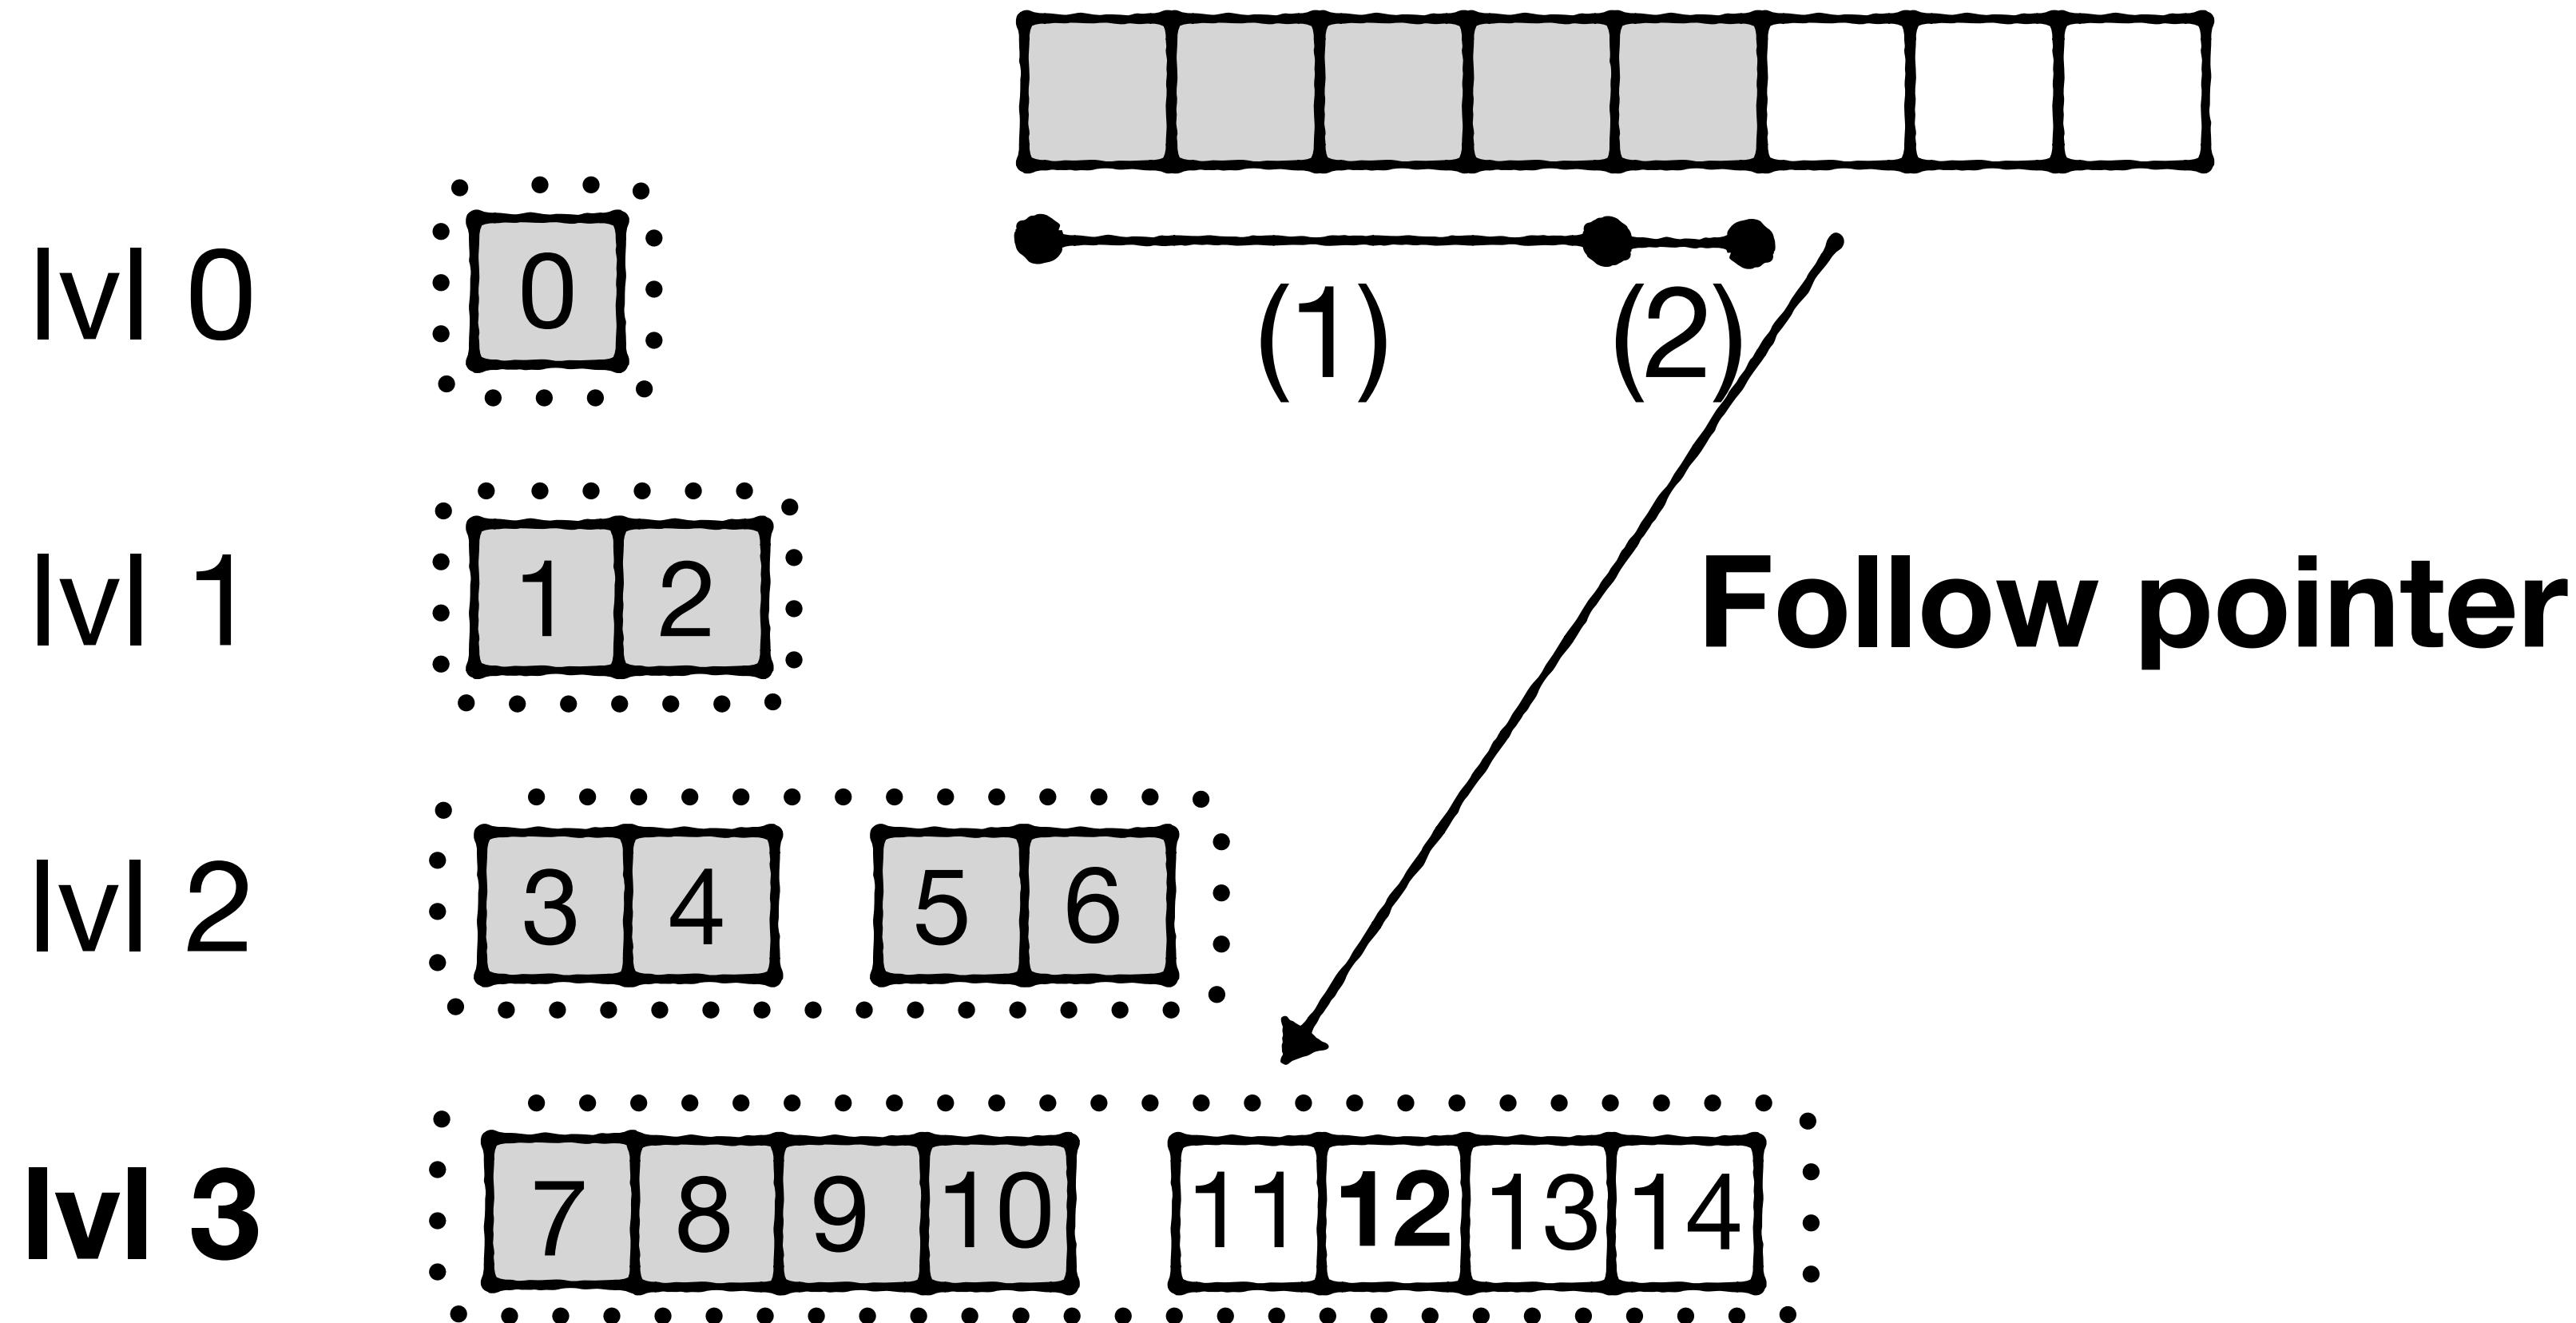


**(1) # blocks to skip in
smaller levels - tricky**

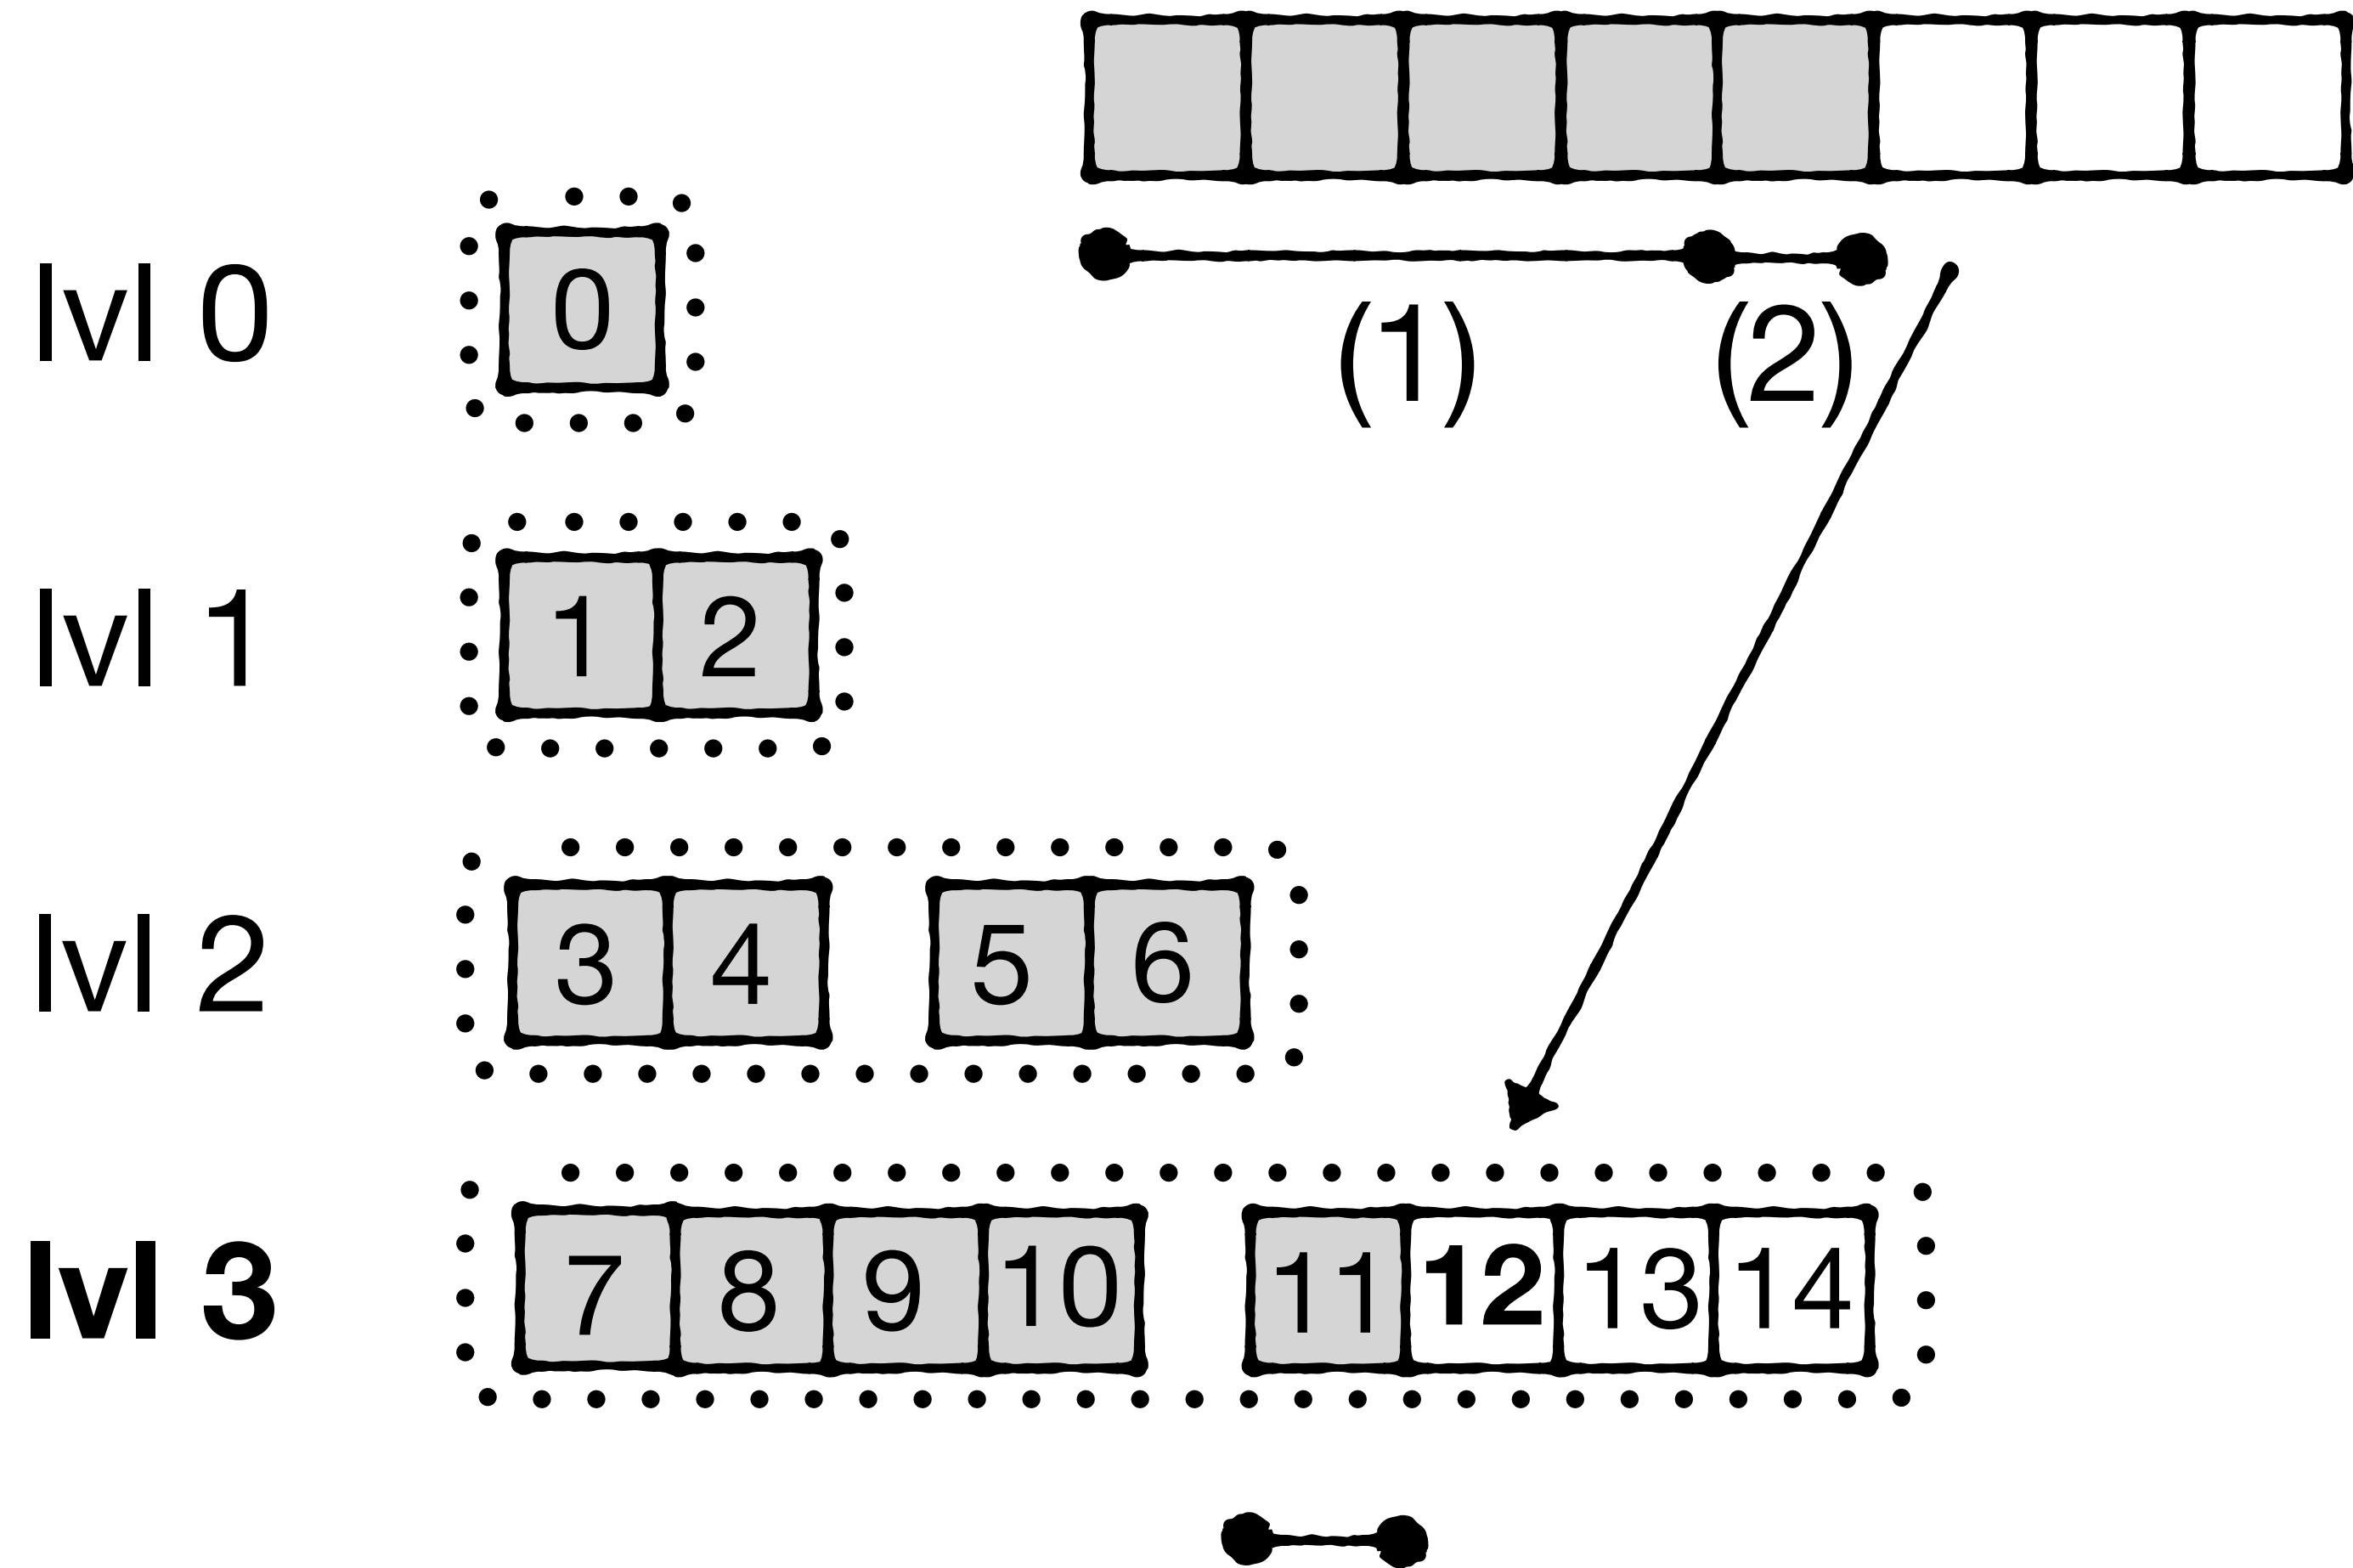
get(12)



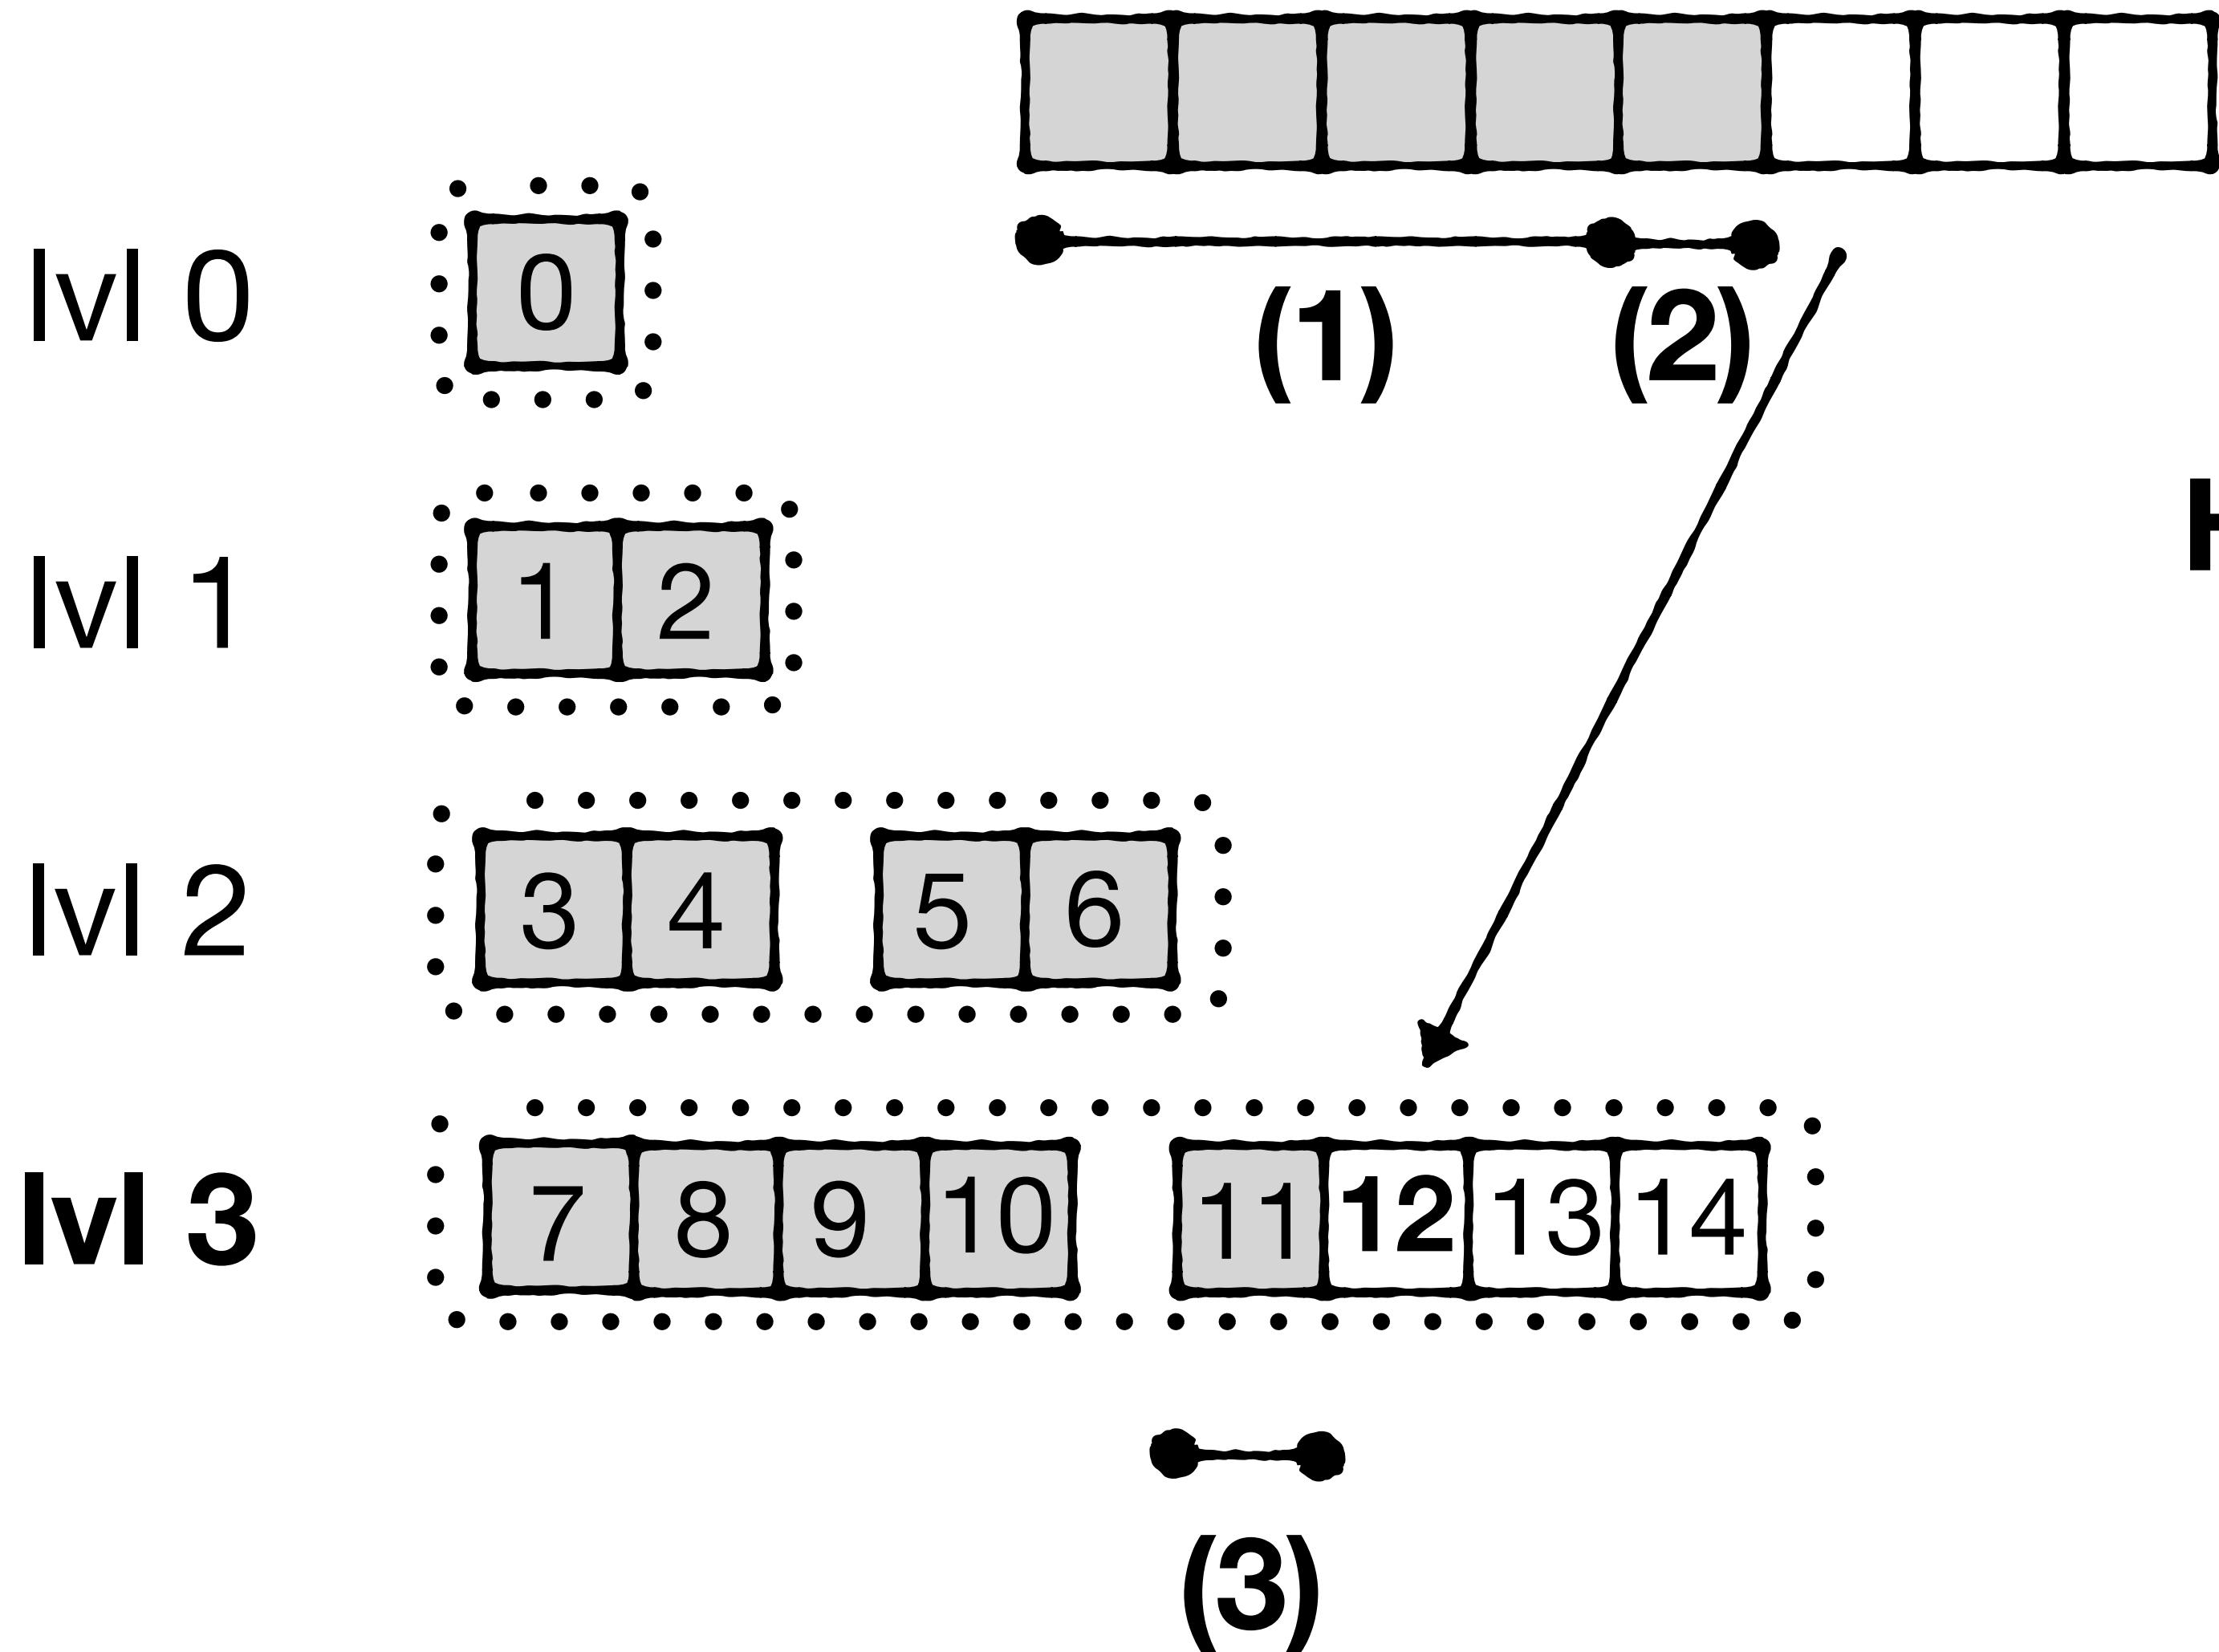
get(12)



get(12)

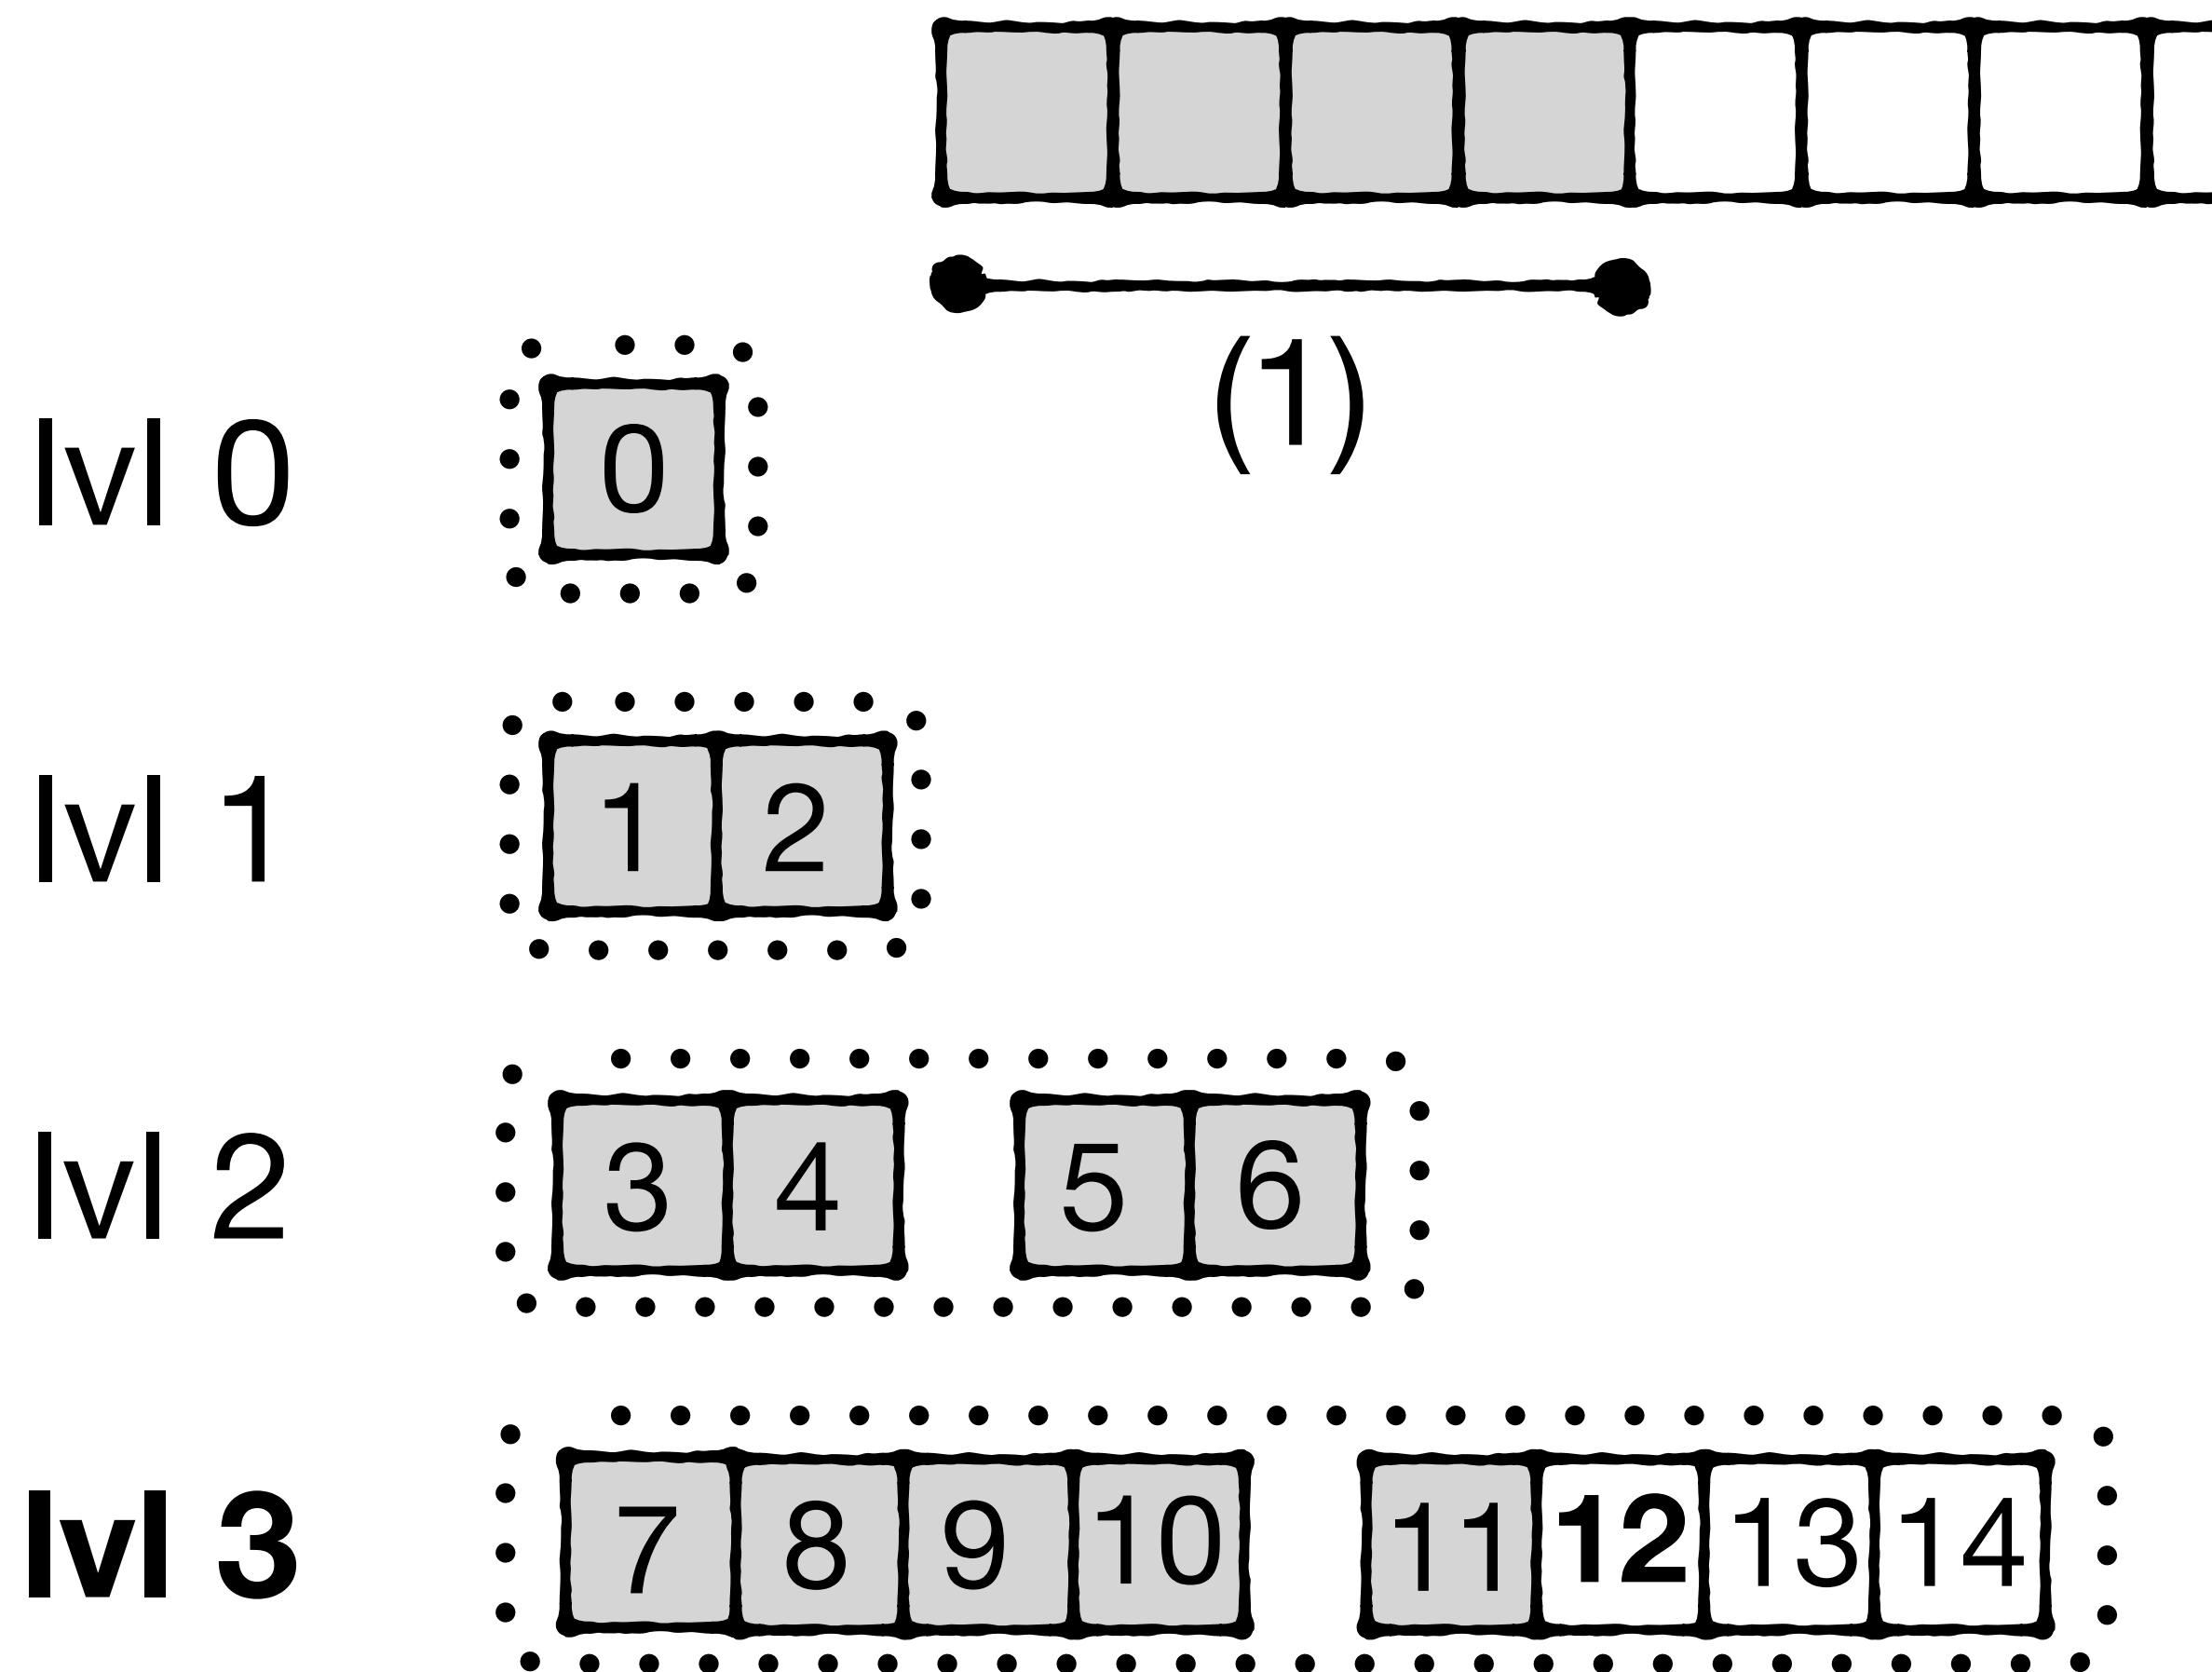


**(3) # slots to skip within
target block**

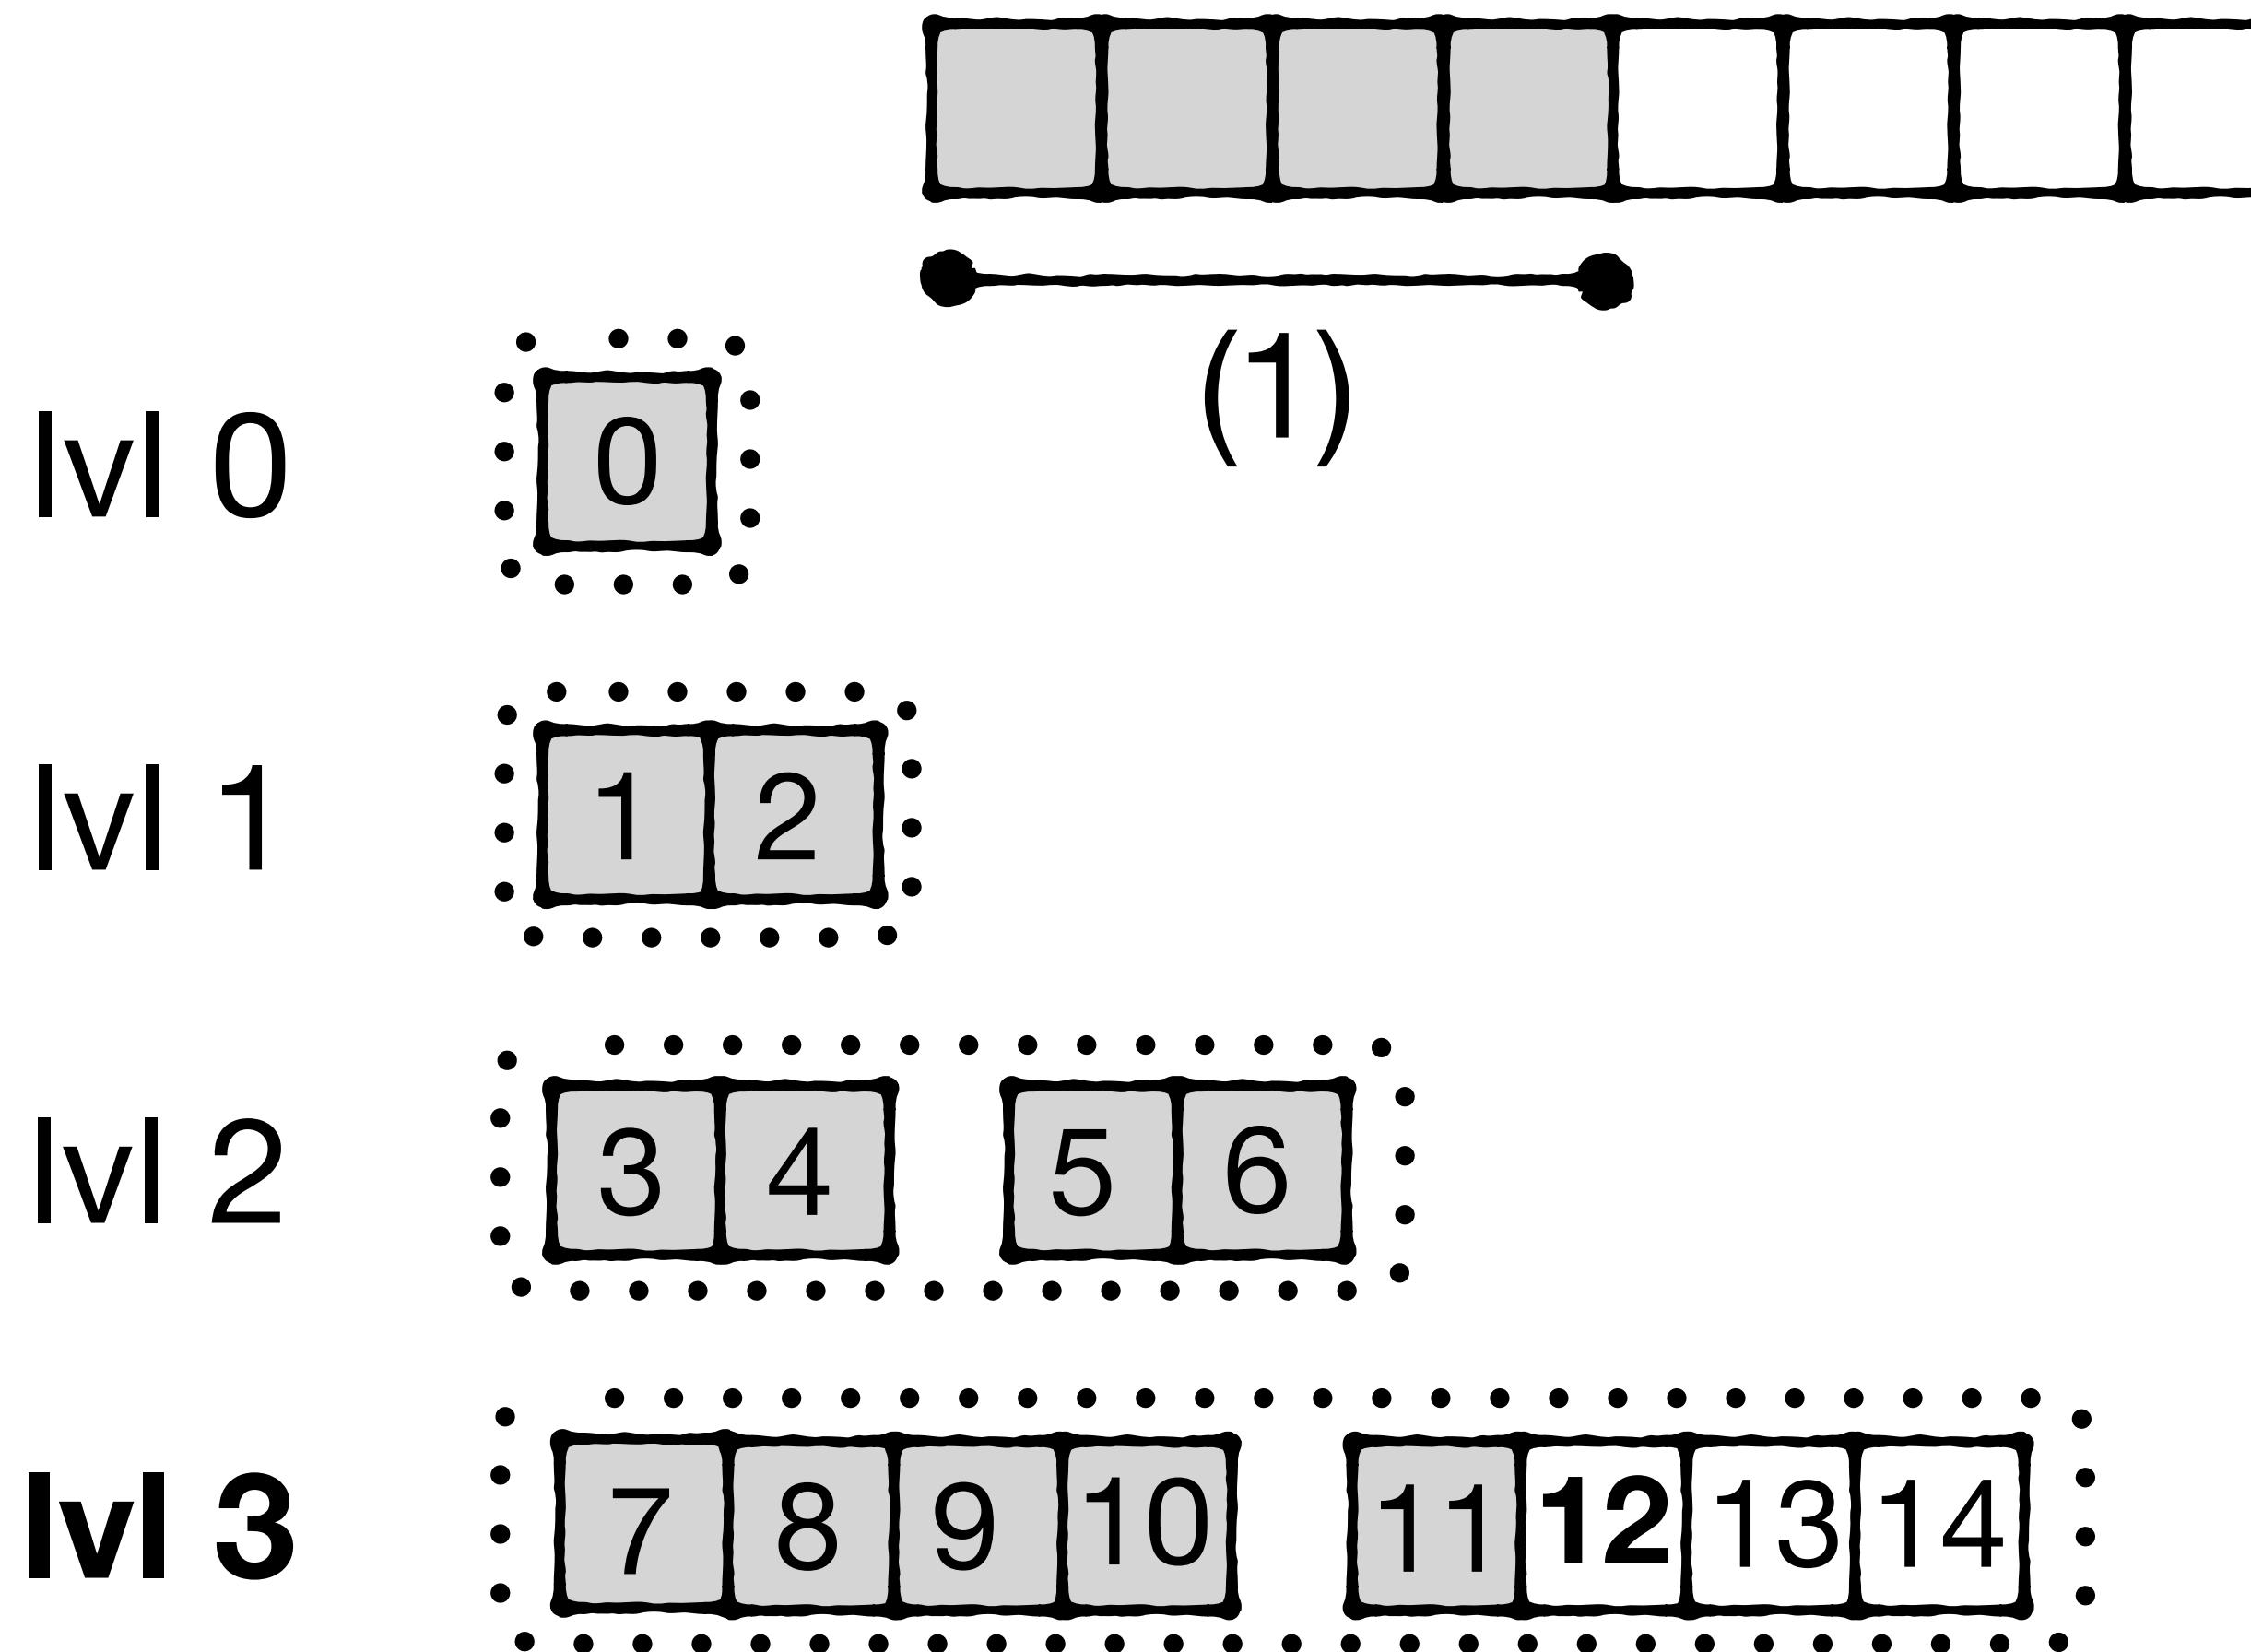


**How to do steps 1 to 3
super fast?**

get(i)

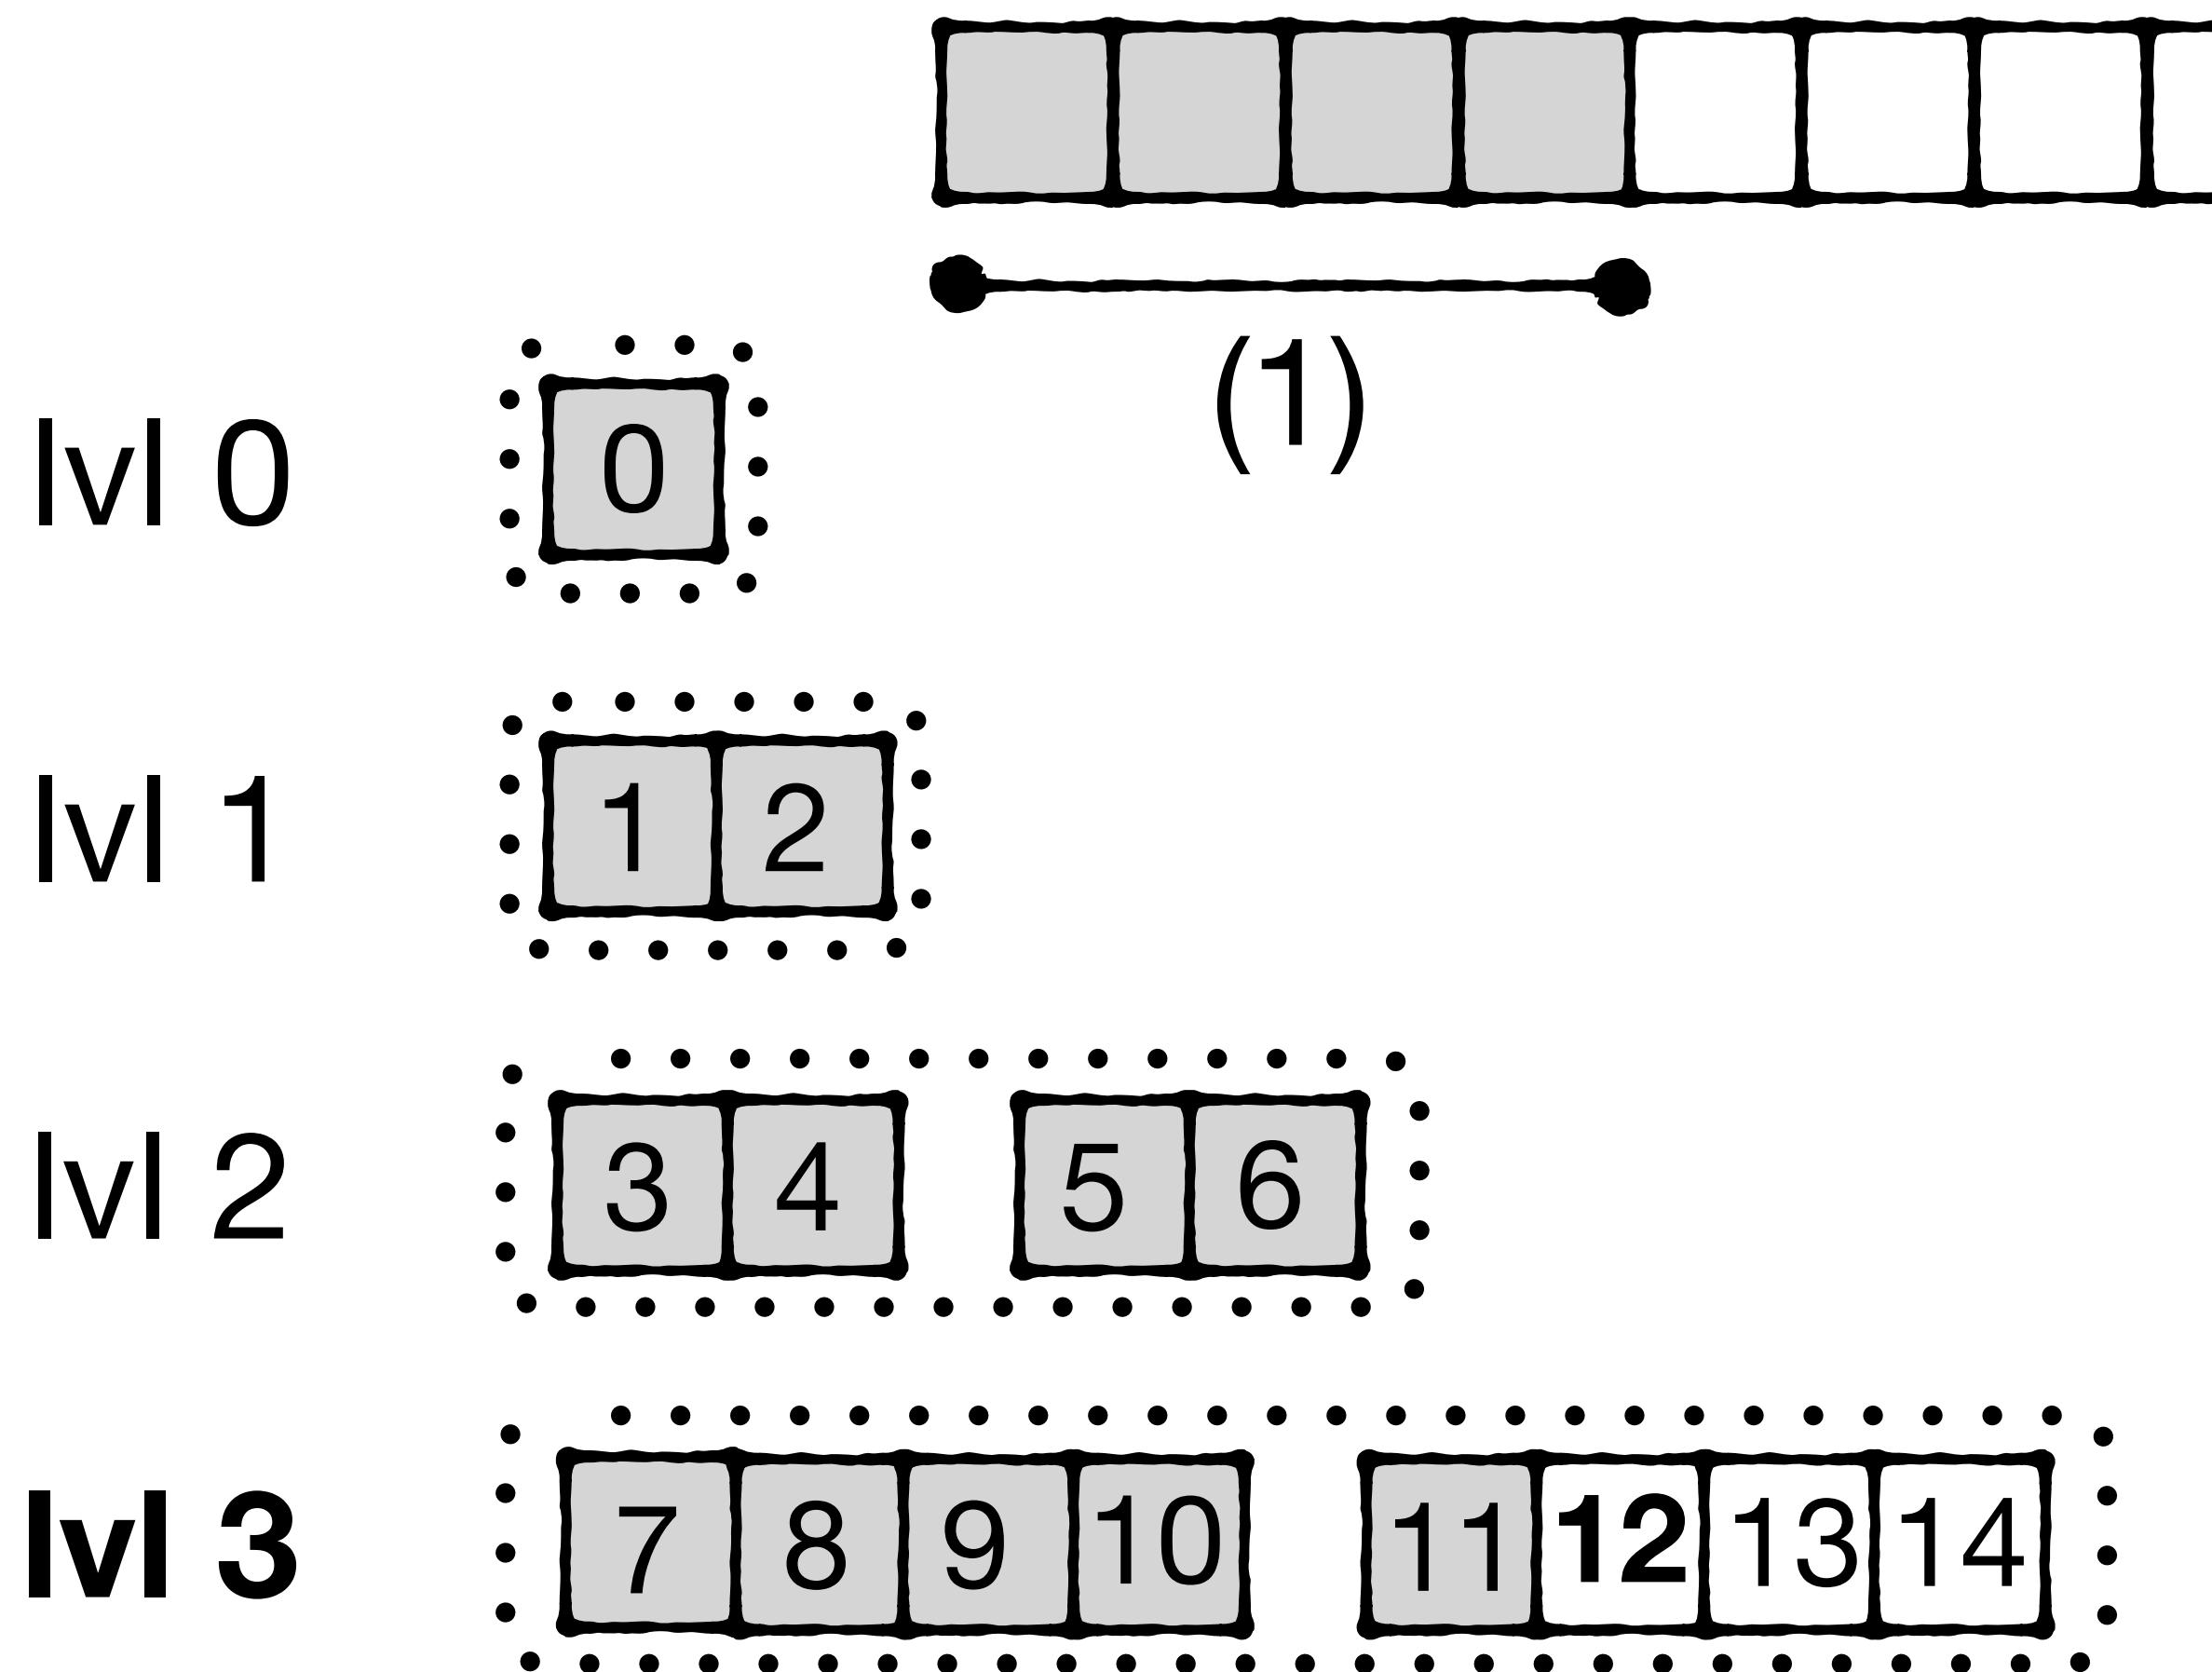


get(i)



Identify target level k

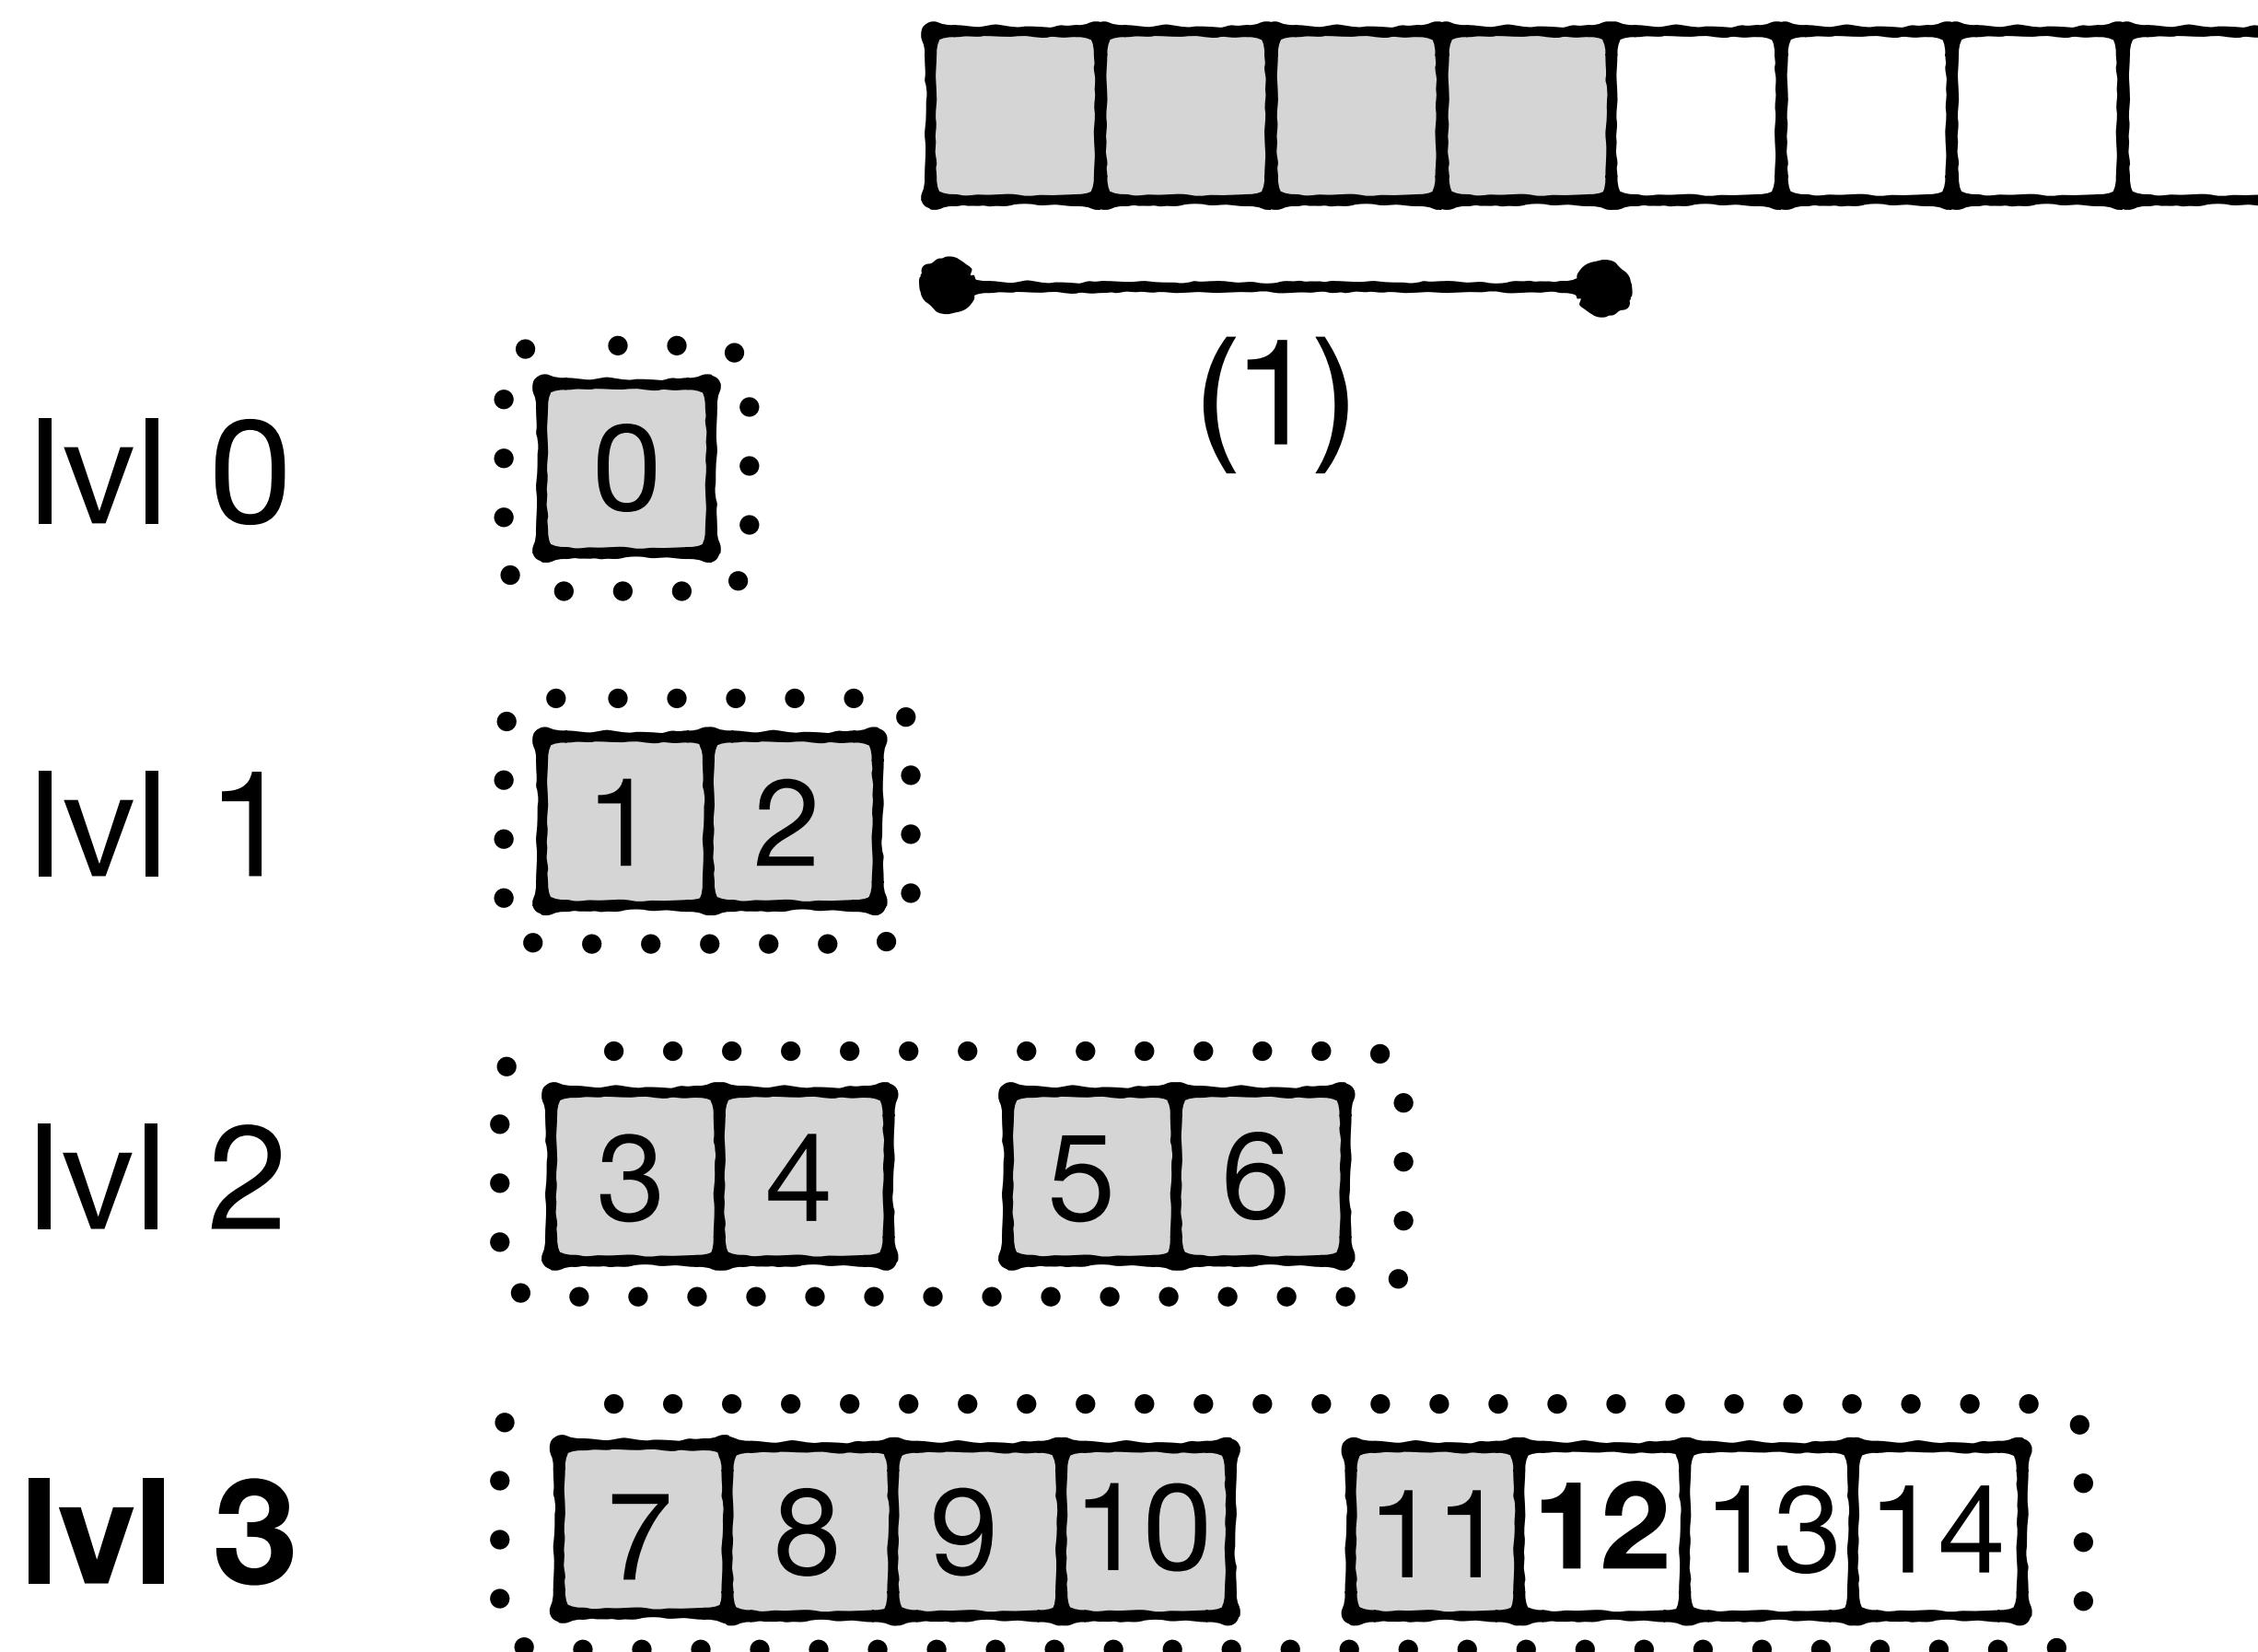
get(i)



Identify target level k

$$k = \lfloor \log_2(i + 1) \rfloor$$

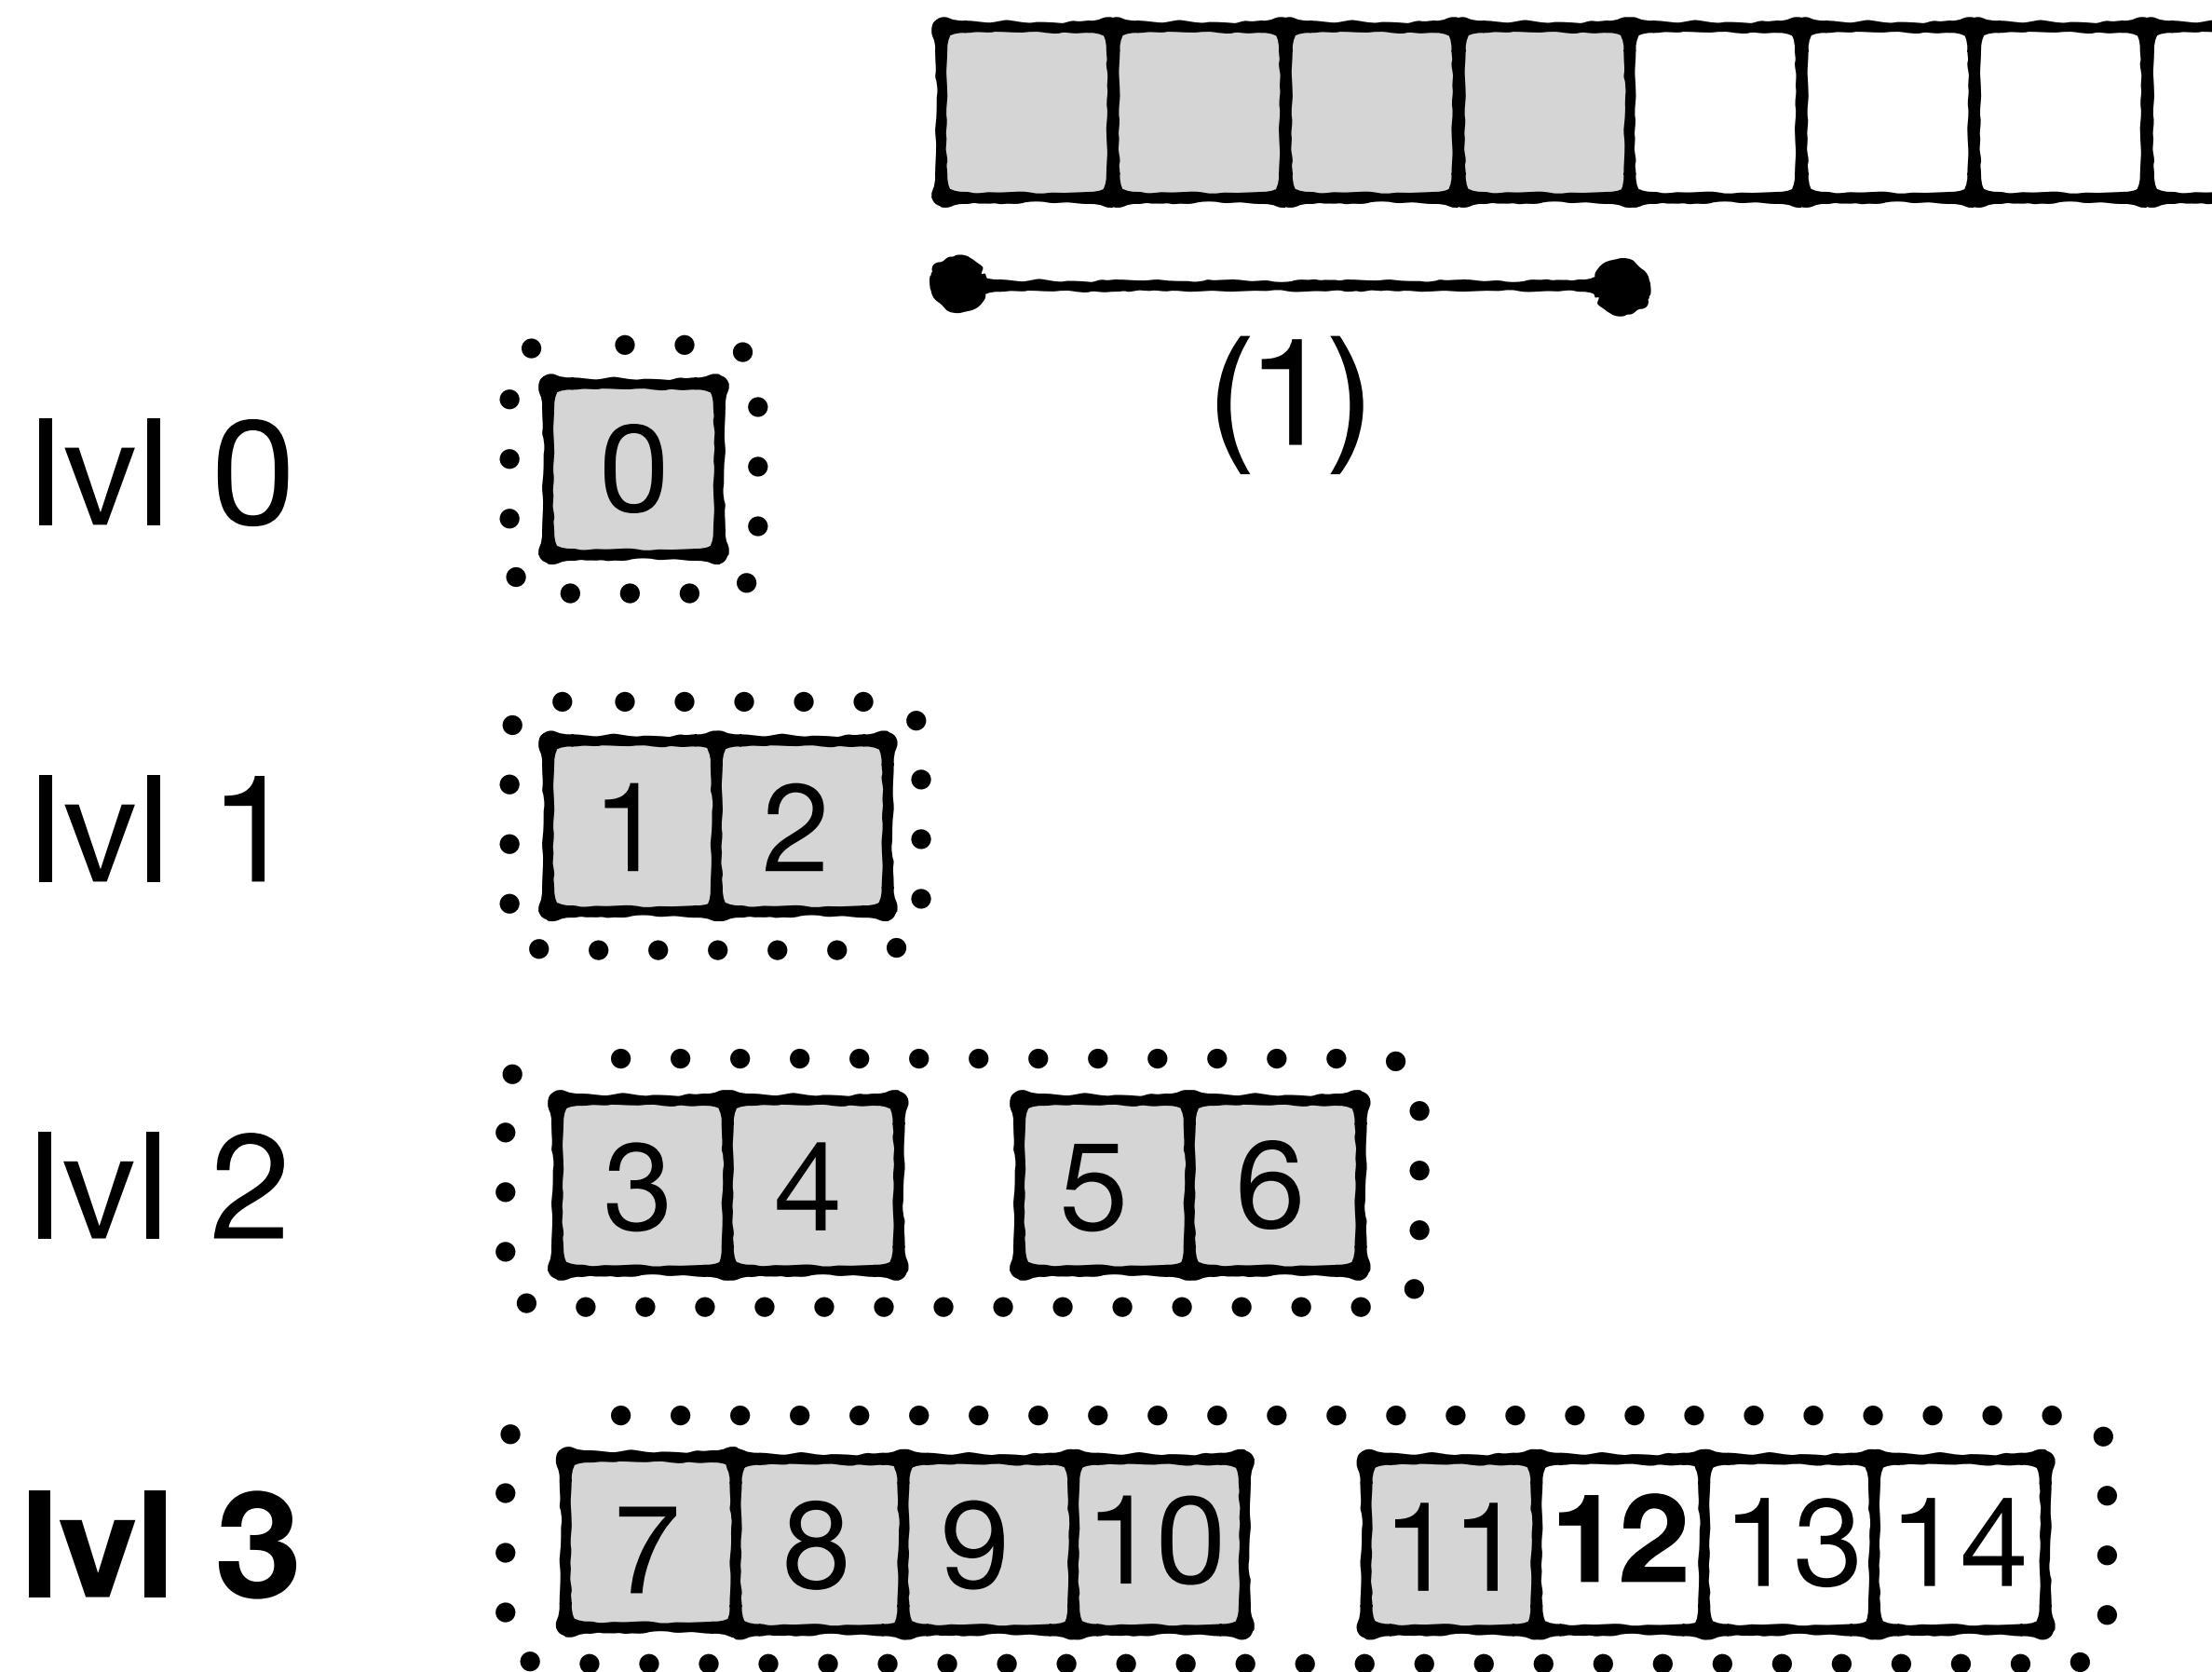
get(12)



Identify target level k

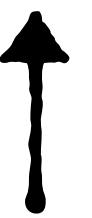
$$k = \lfloor \log_2(12 + 1) \rfloor = 3$$

get(i)



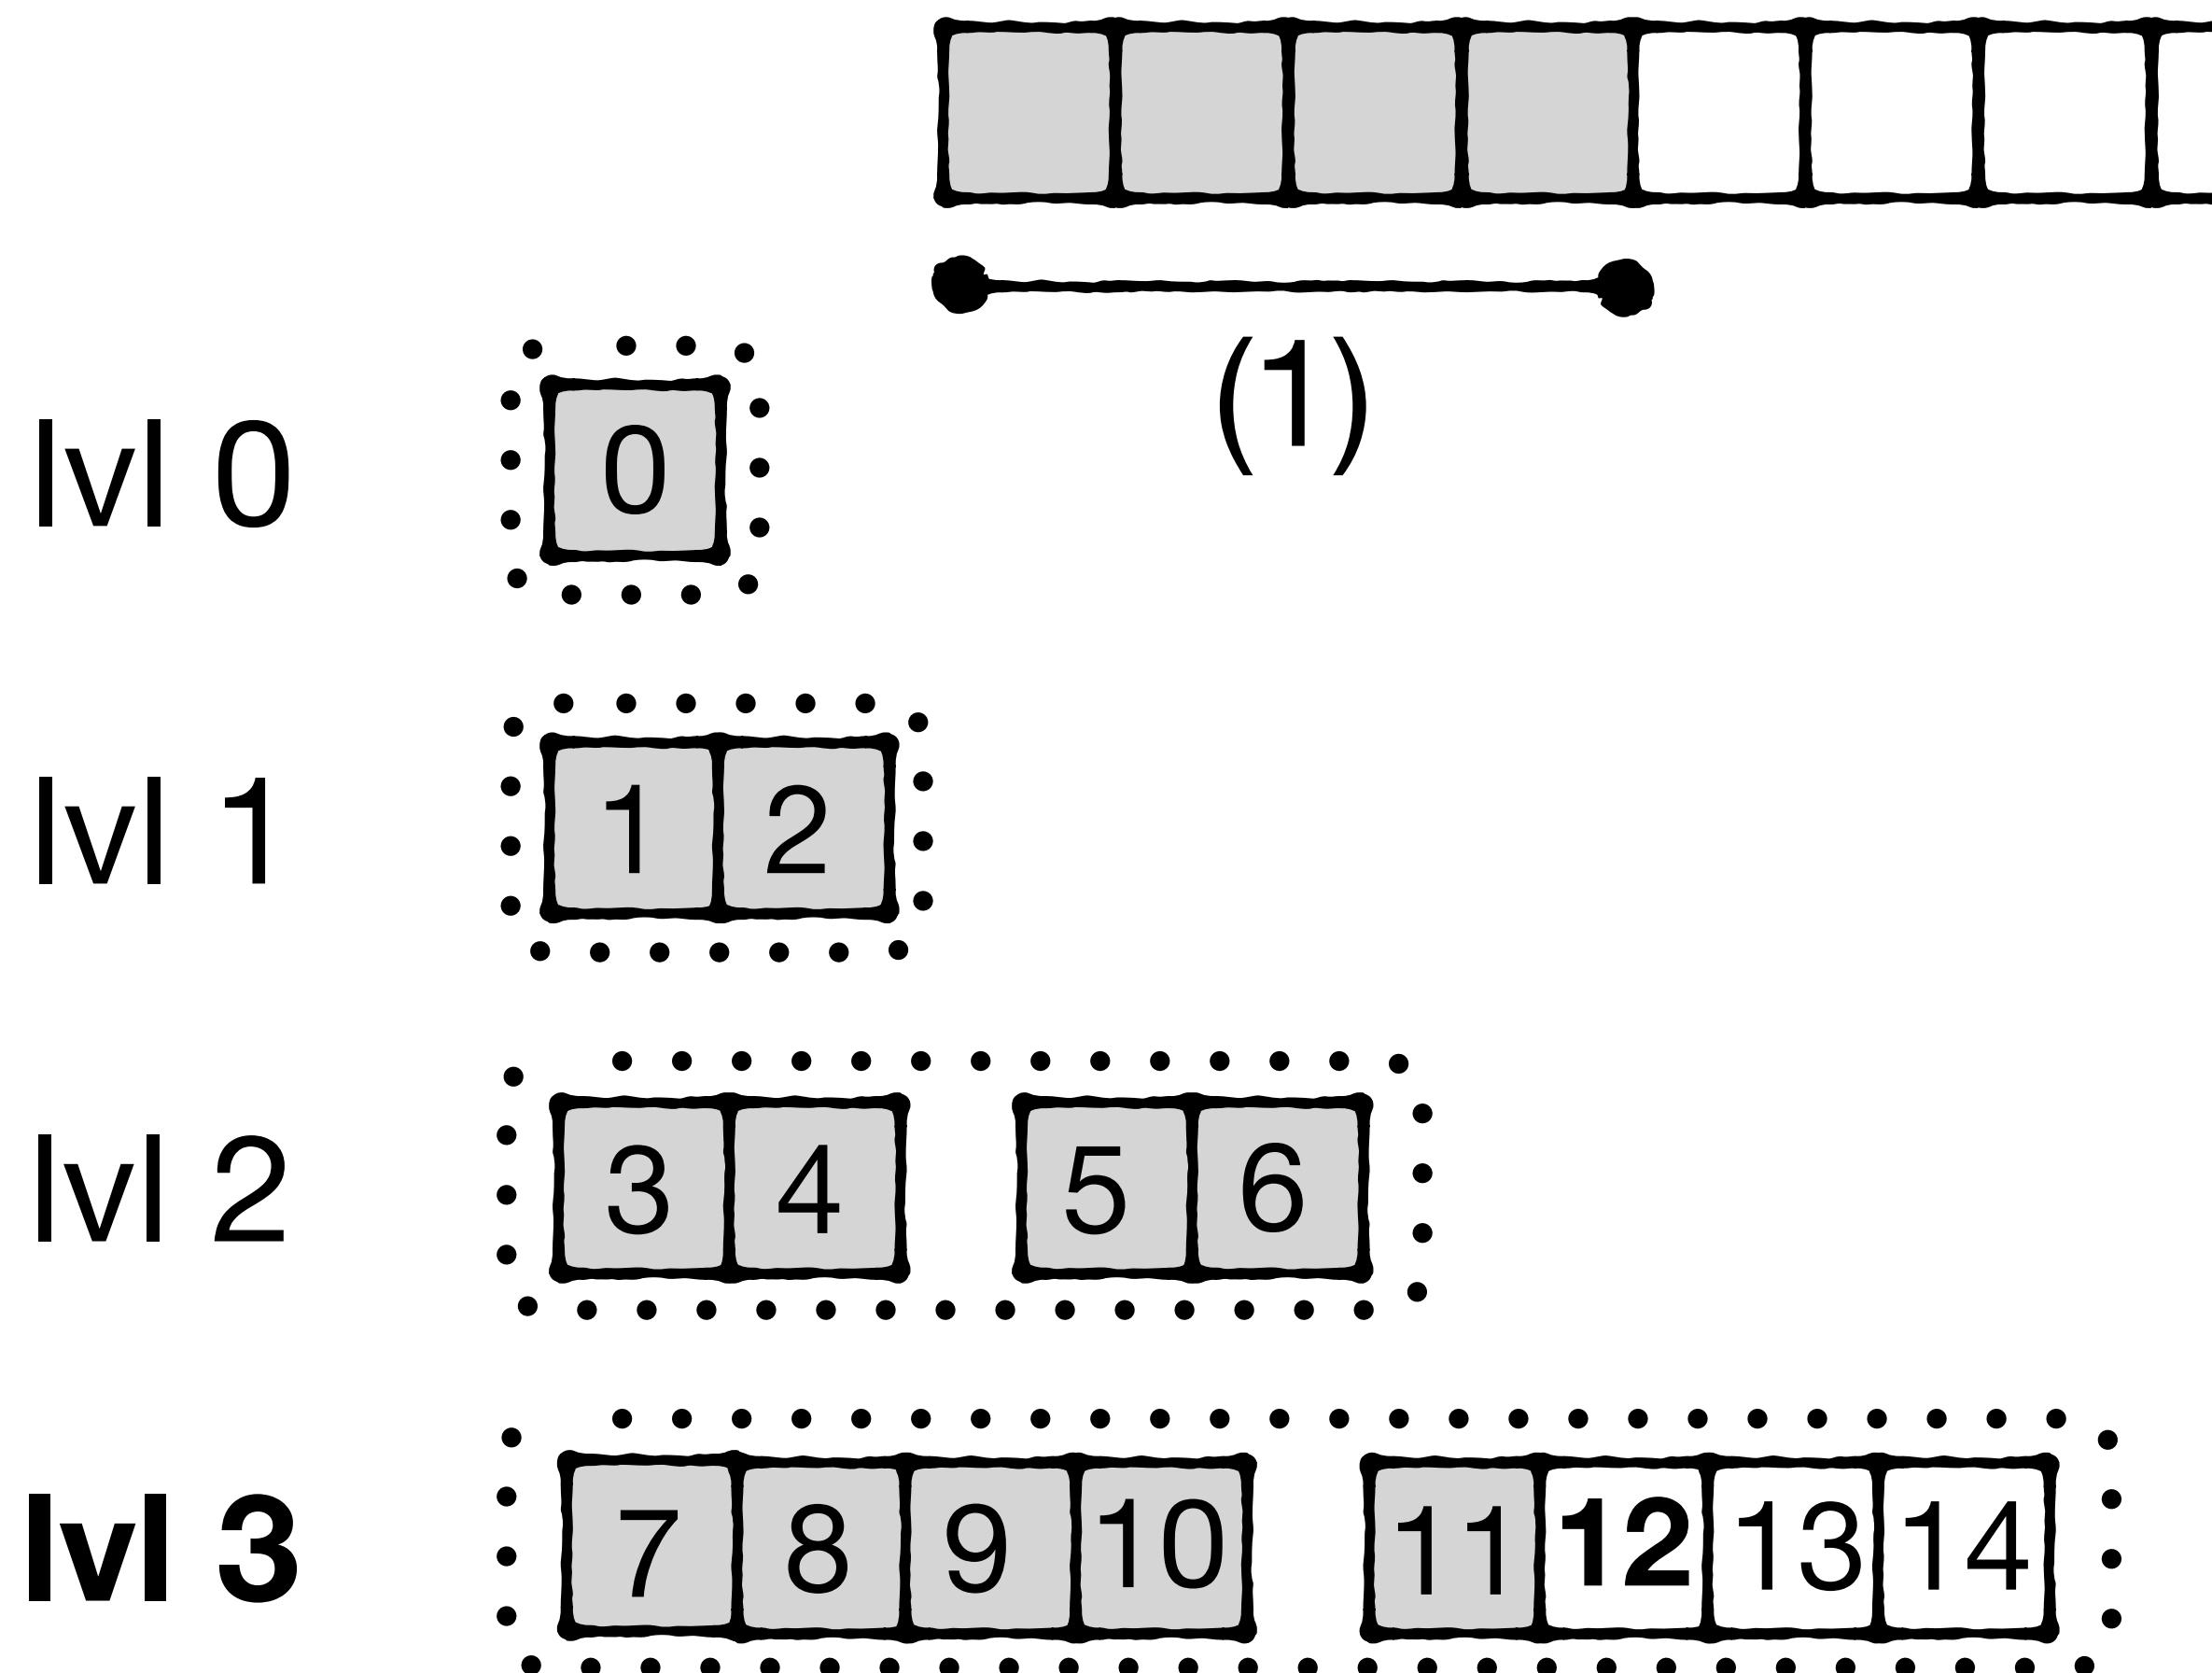
Identify target level k

$$k = \lfloor \log_2(i + 1) \rfloor$$



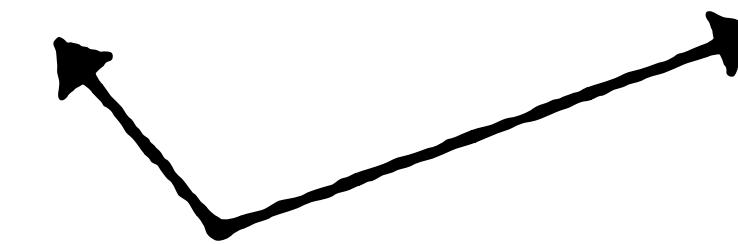
slow

get(i)



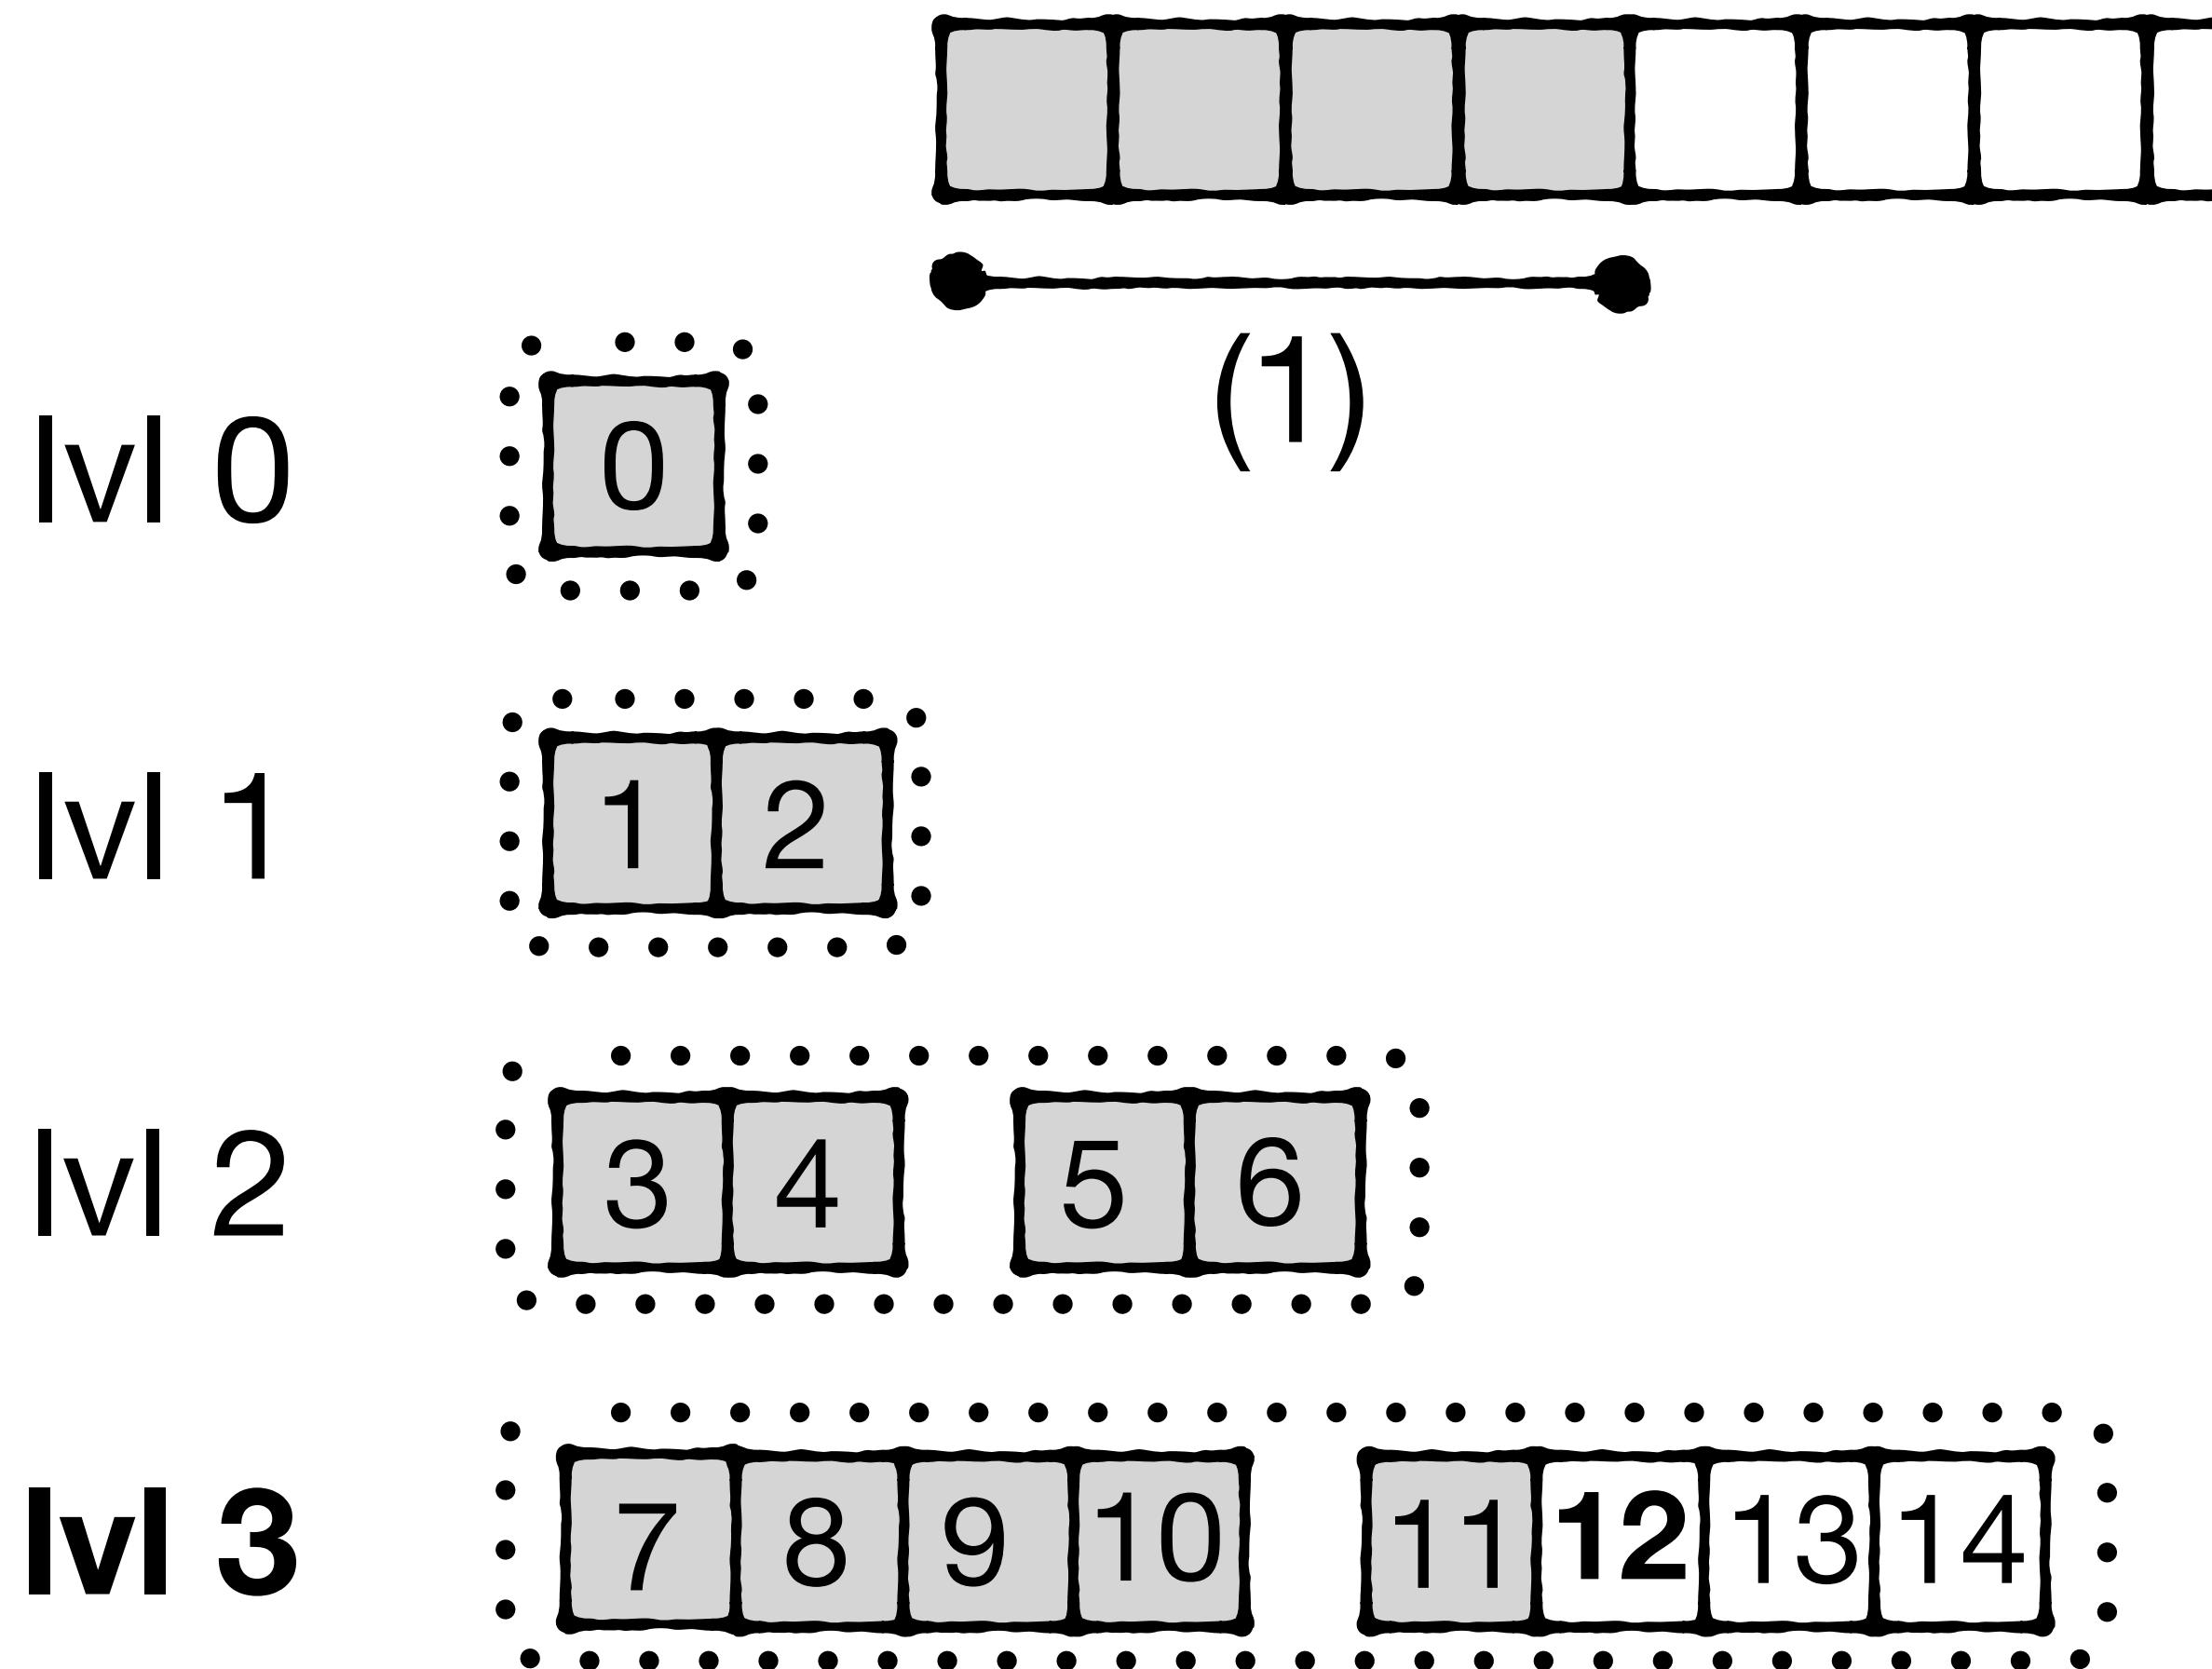
Identify target level k

$$k = \lfloor \log_2(i + 1) \rfloor$$



Type casting - also slow

get(i)

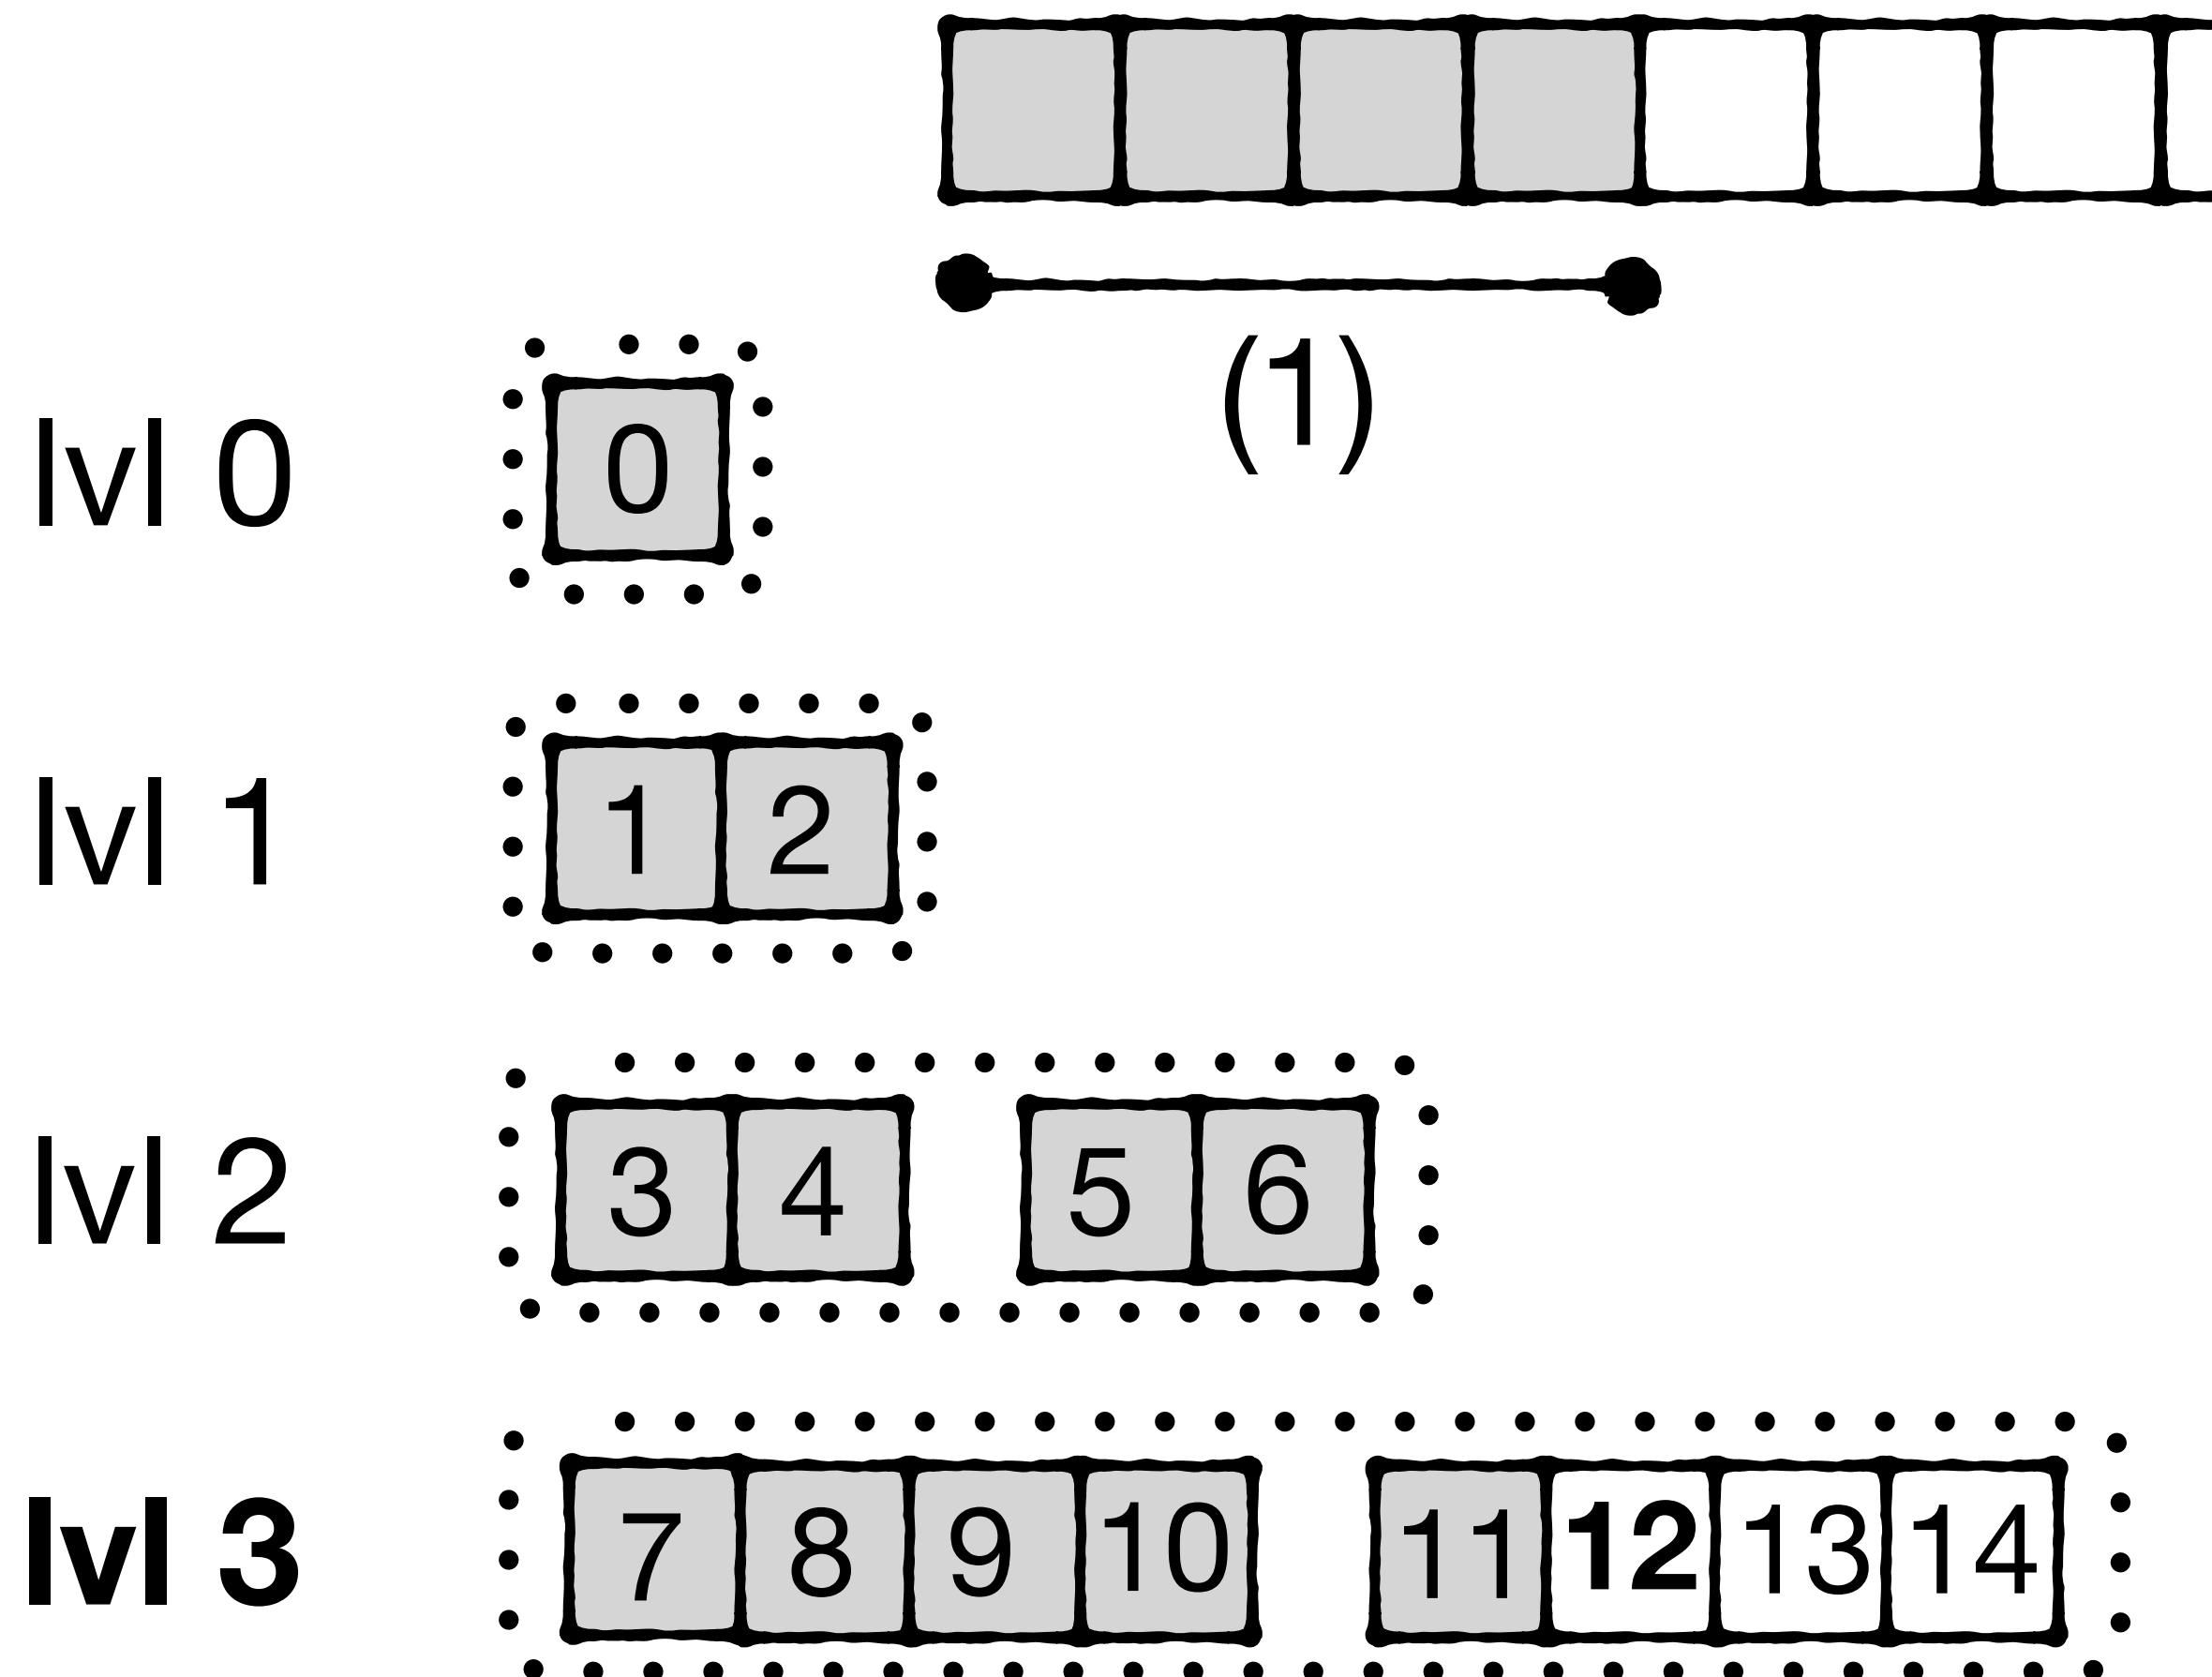


Identify target level k

$$k = \lfloor \log_2(i + 1) \rfloor$$

Insight?

get(i)

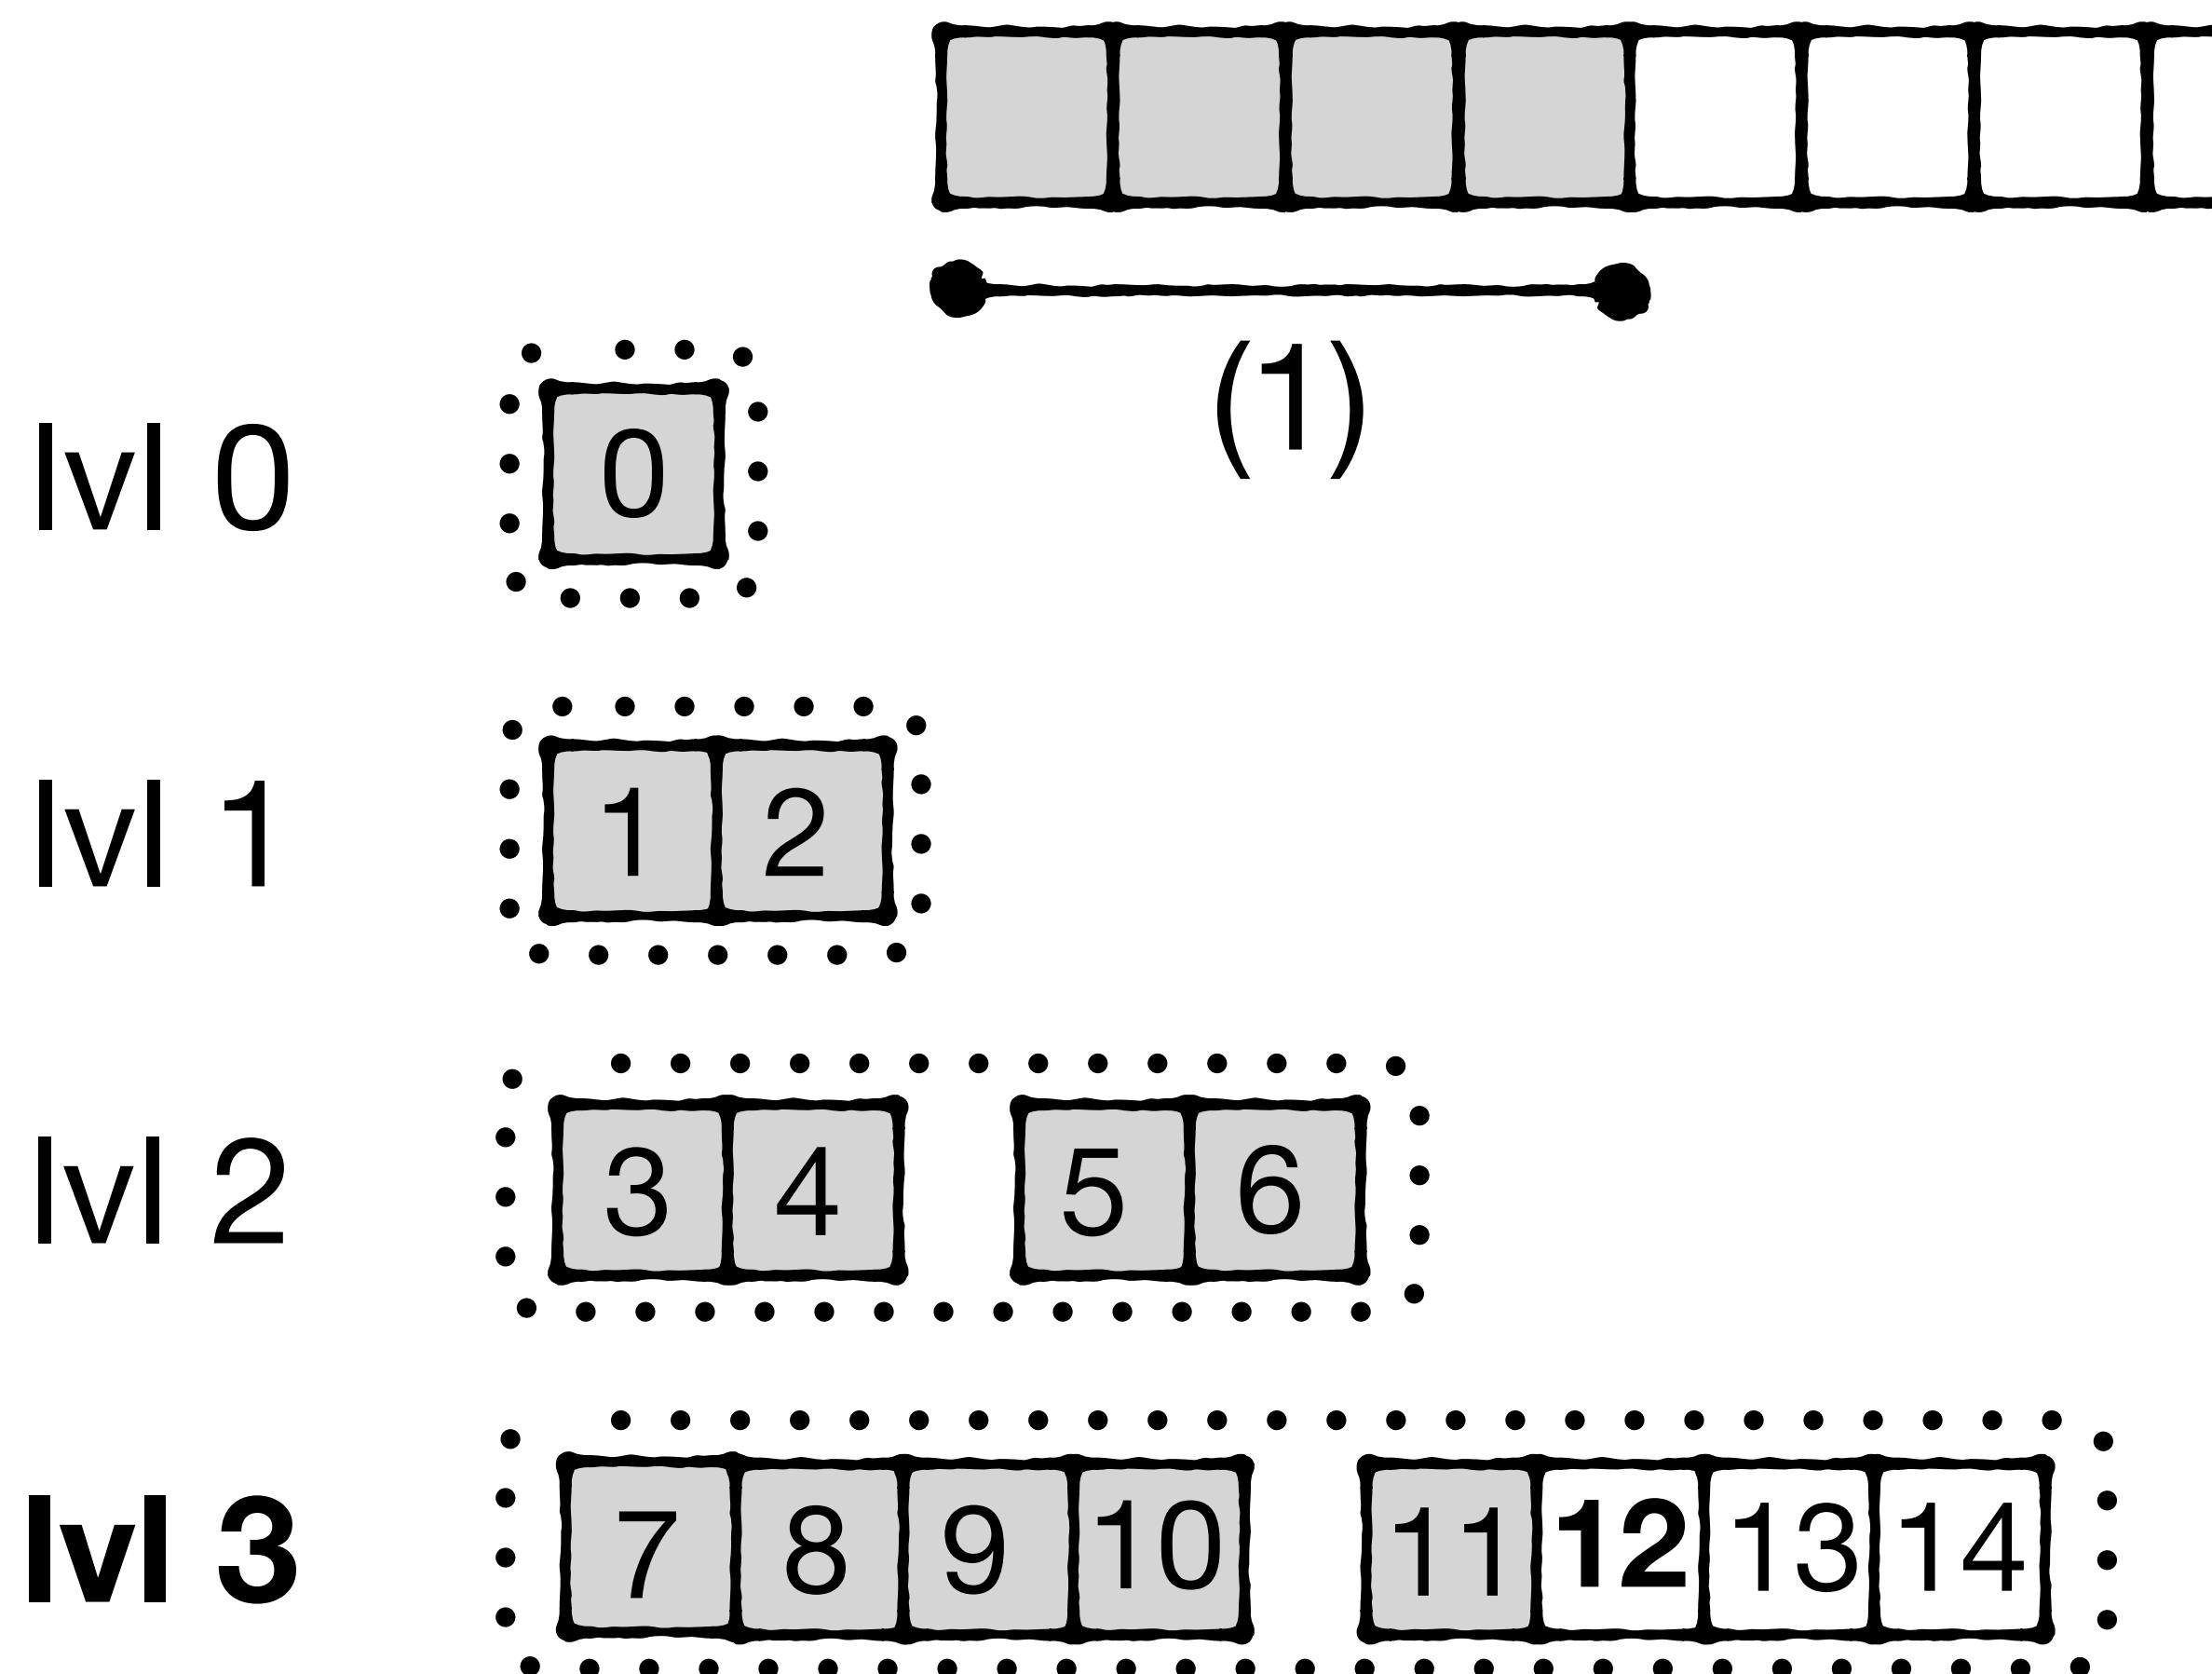


Identify target level k

$$k = \lfloor \log_2(i + 1) \rfloor$$

Insight: \log_2 amounts to finding index of most significant digit

get(i)

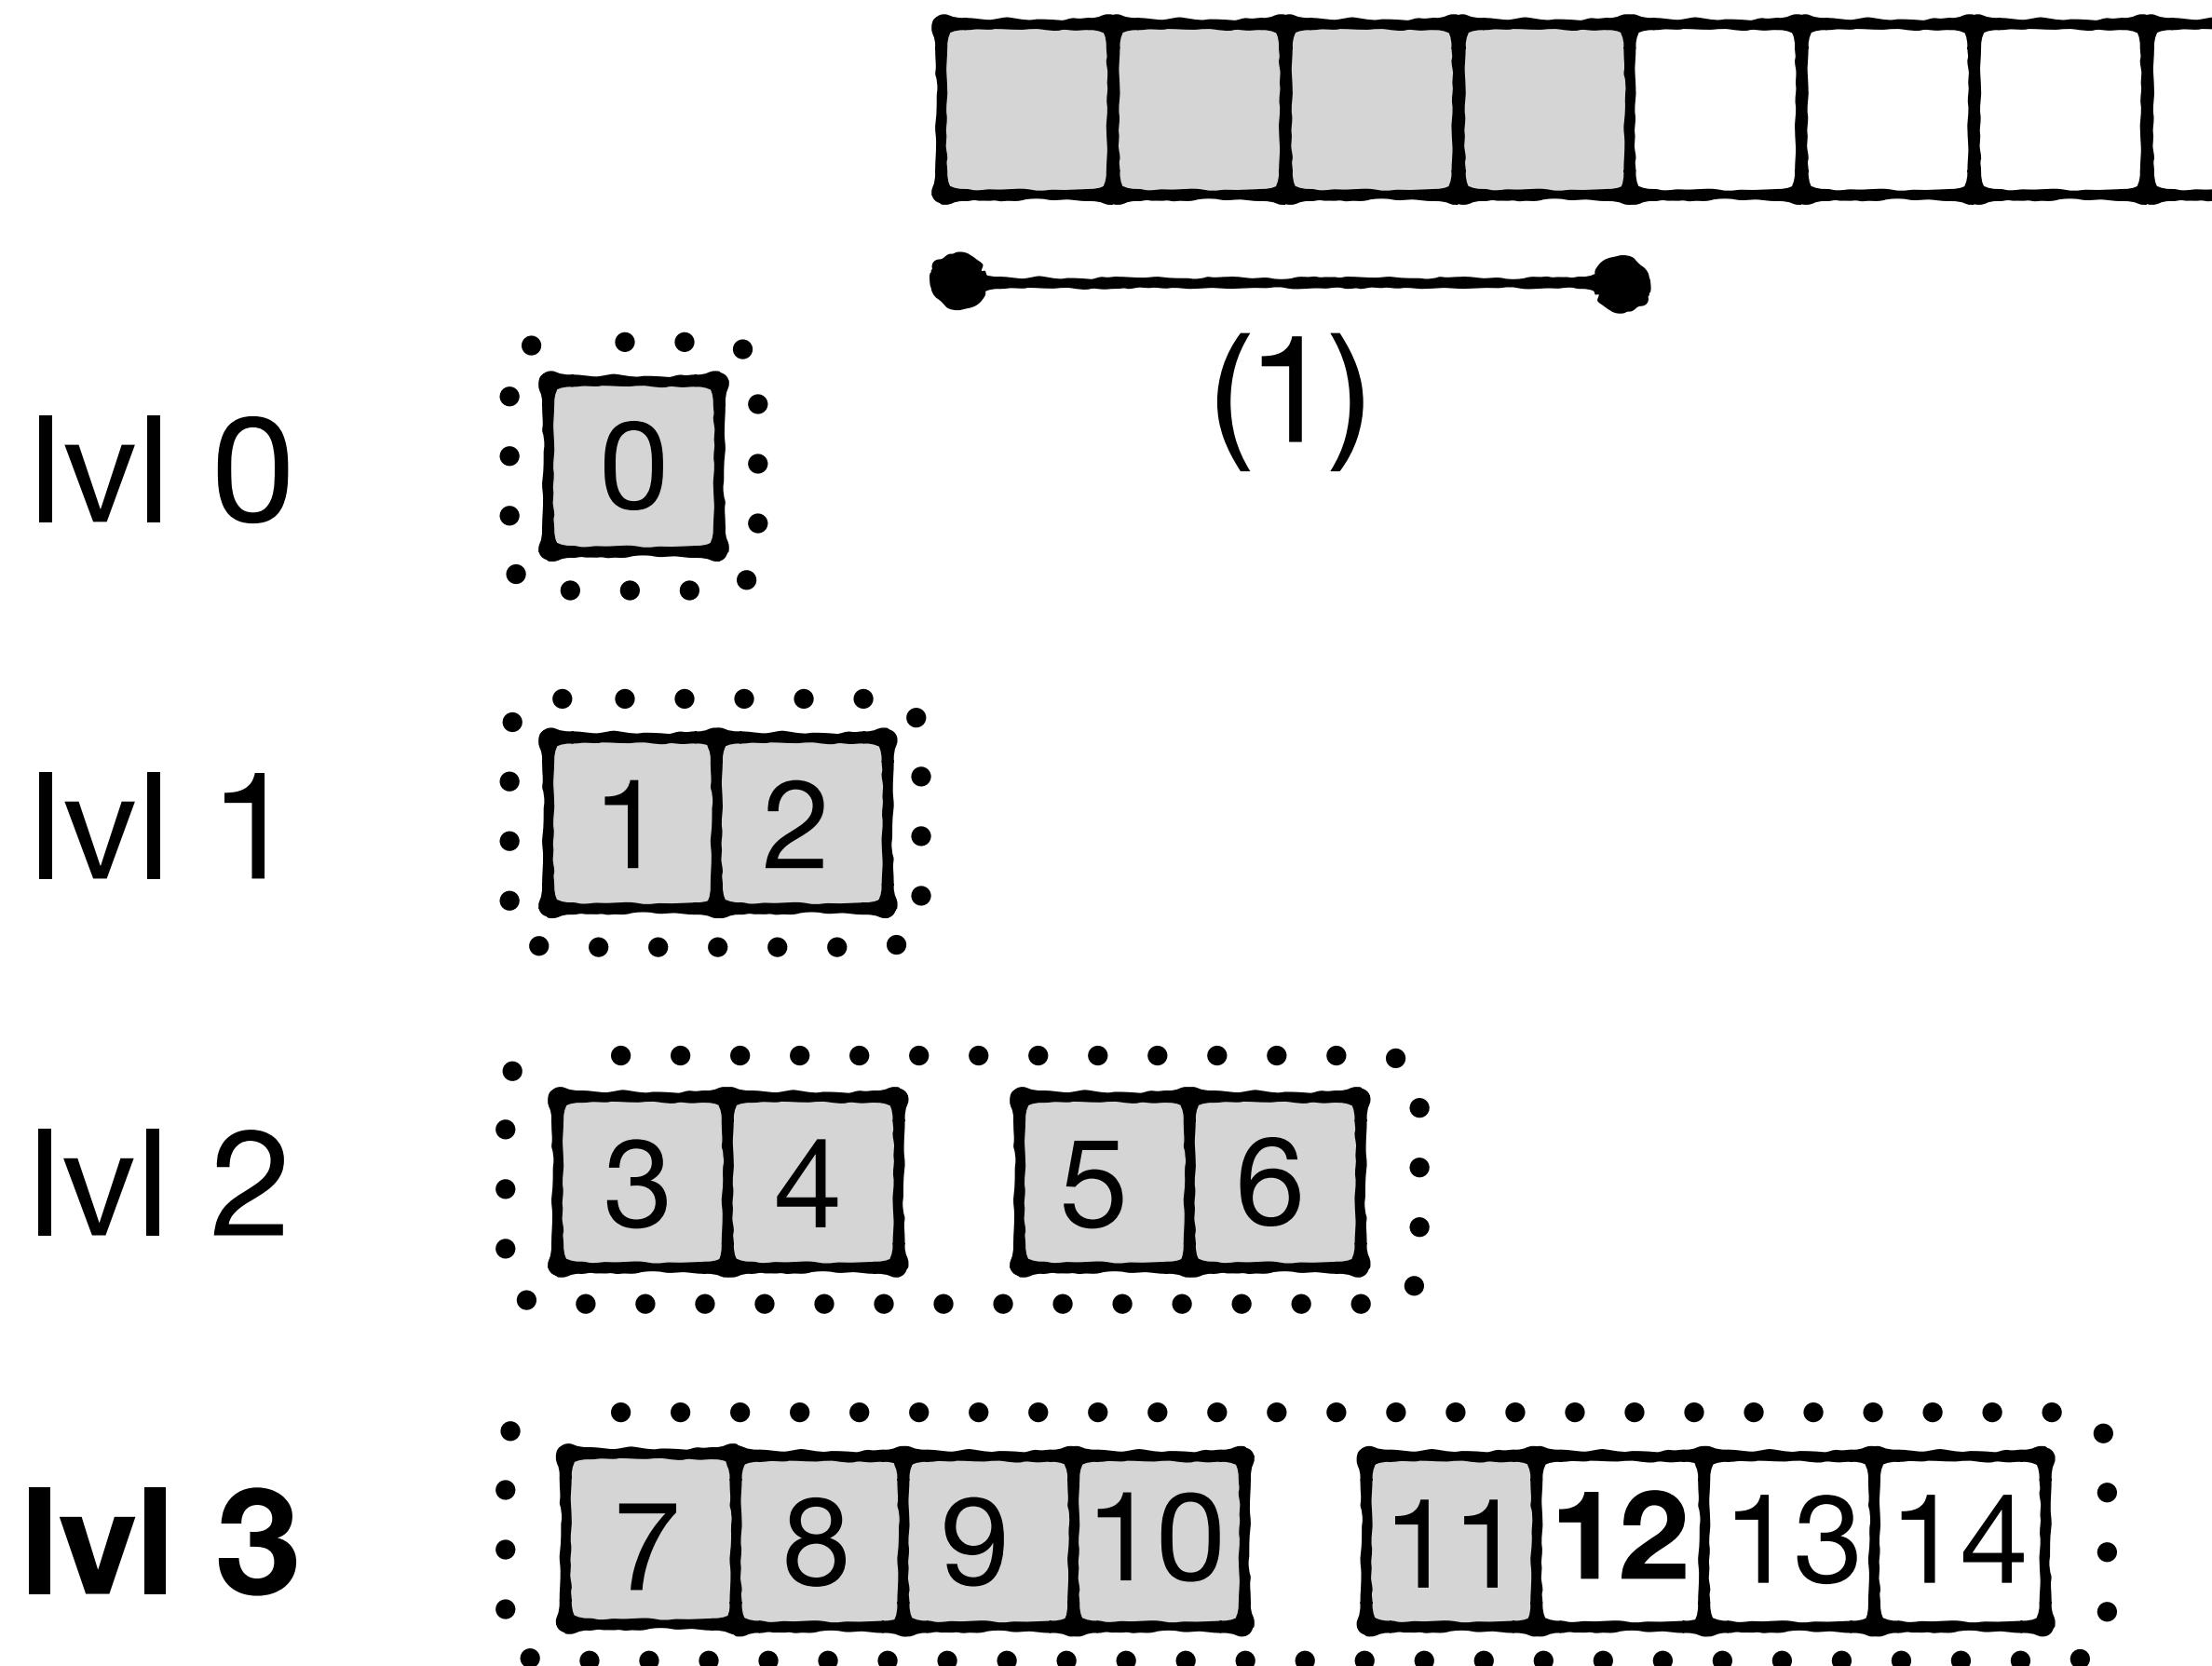


Identify target level k

$$k = \lfloor \log_2(i + 1) \rfloor$$

$$= \text{sizeof}(i) - 1 - \text{clz}(i+1)$$

get(i)



Identify target level k

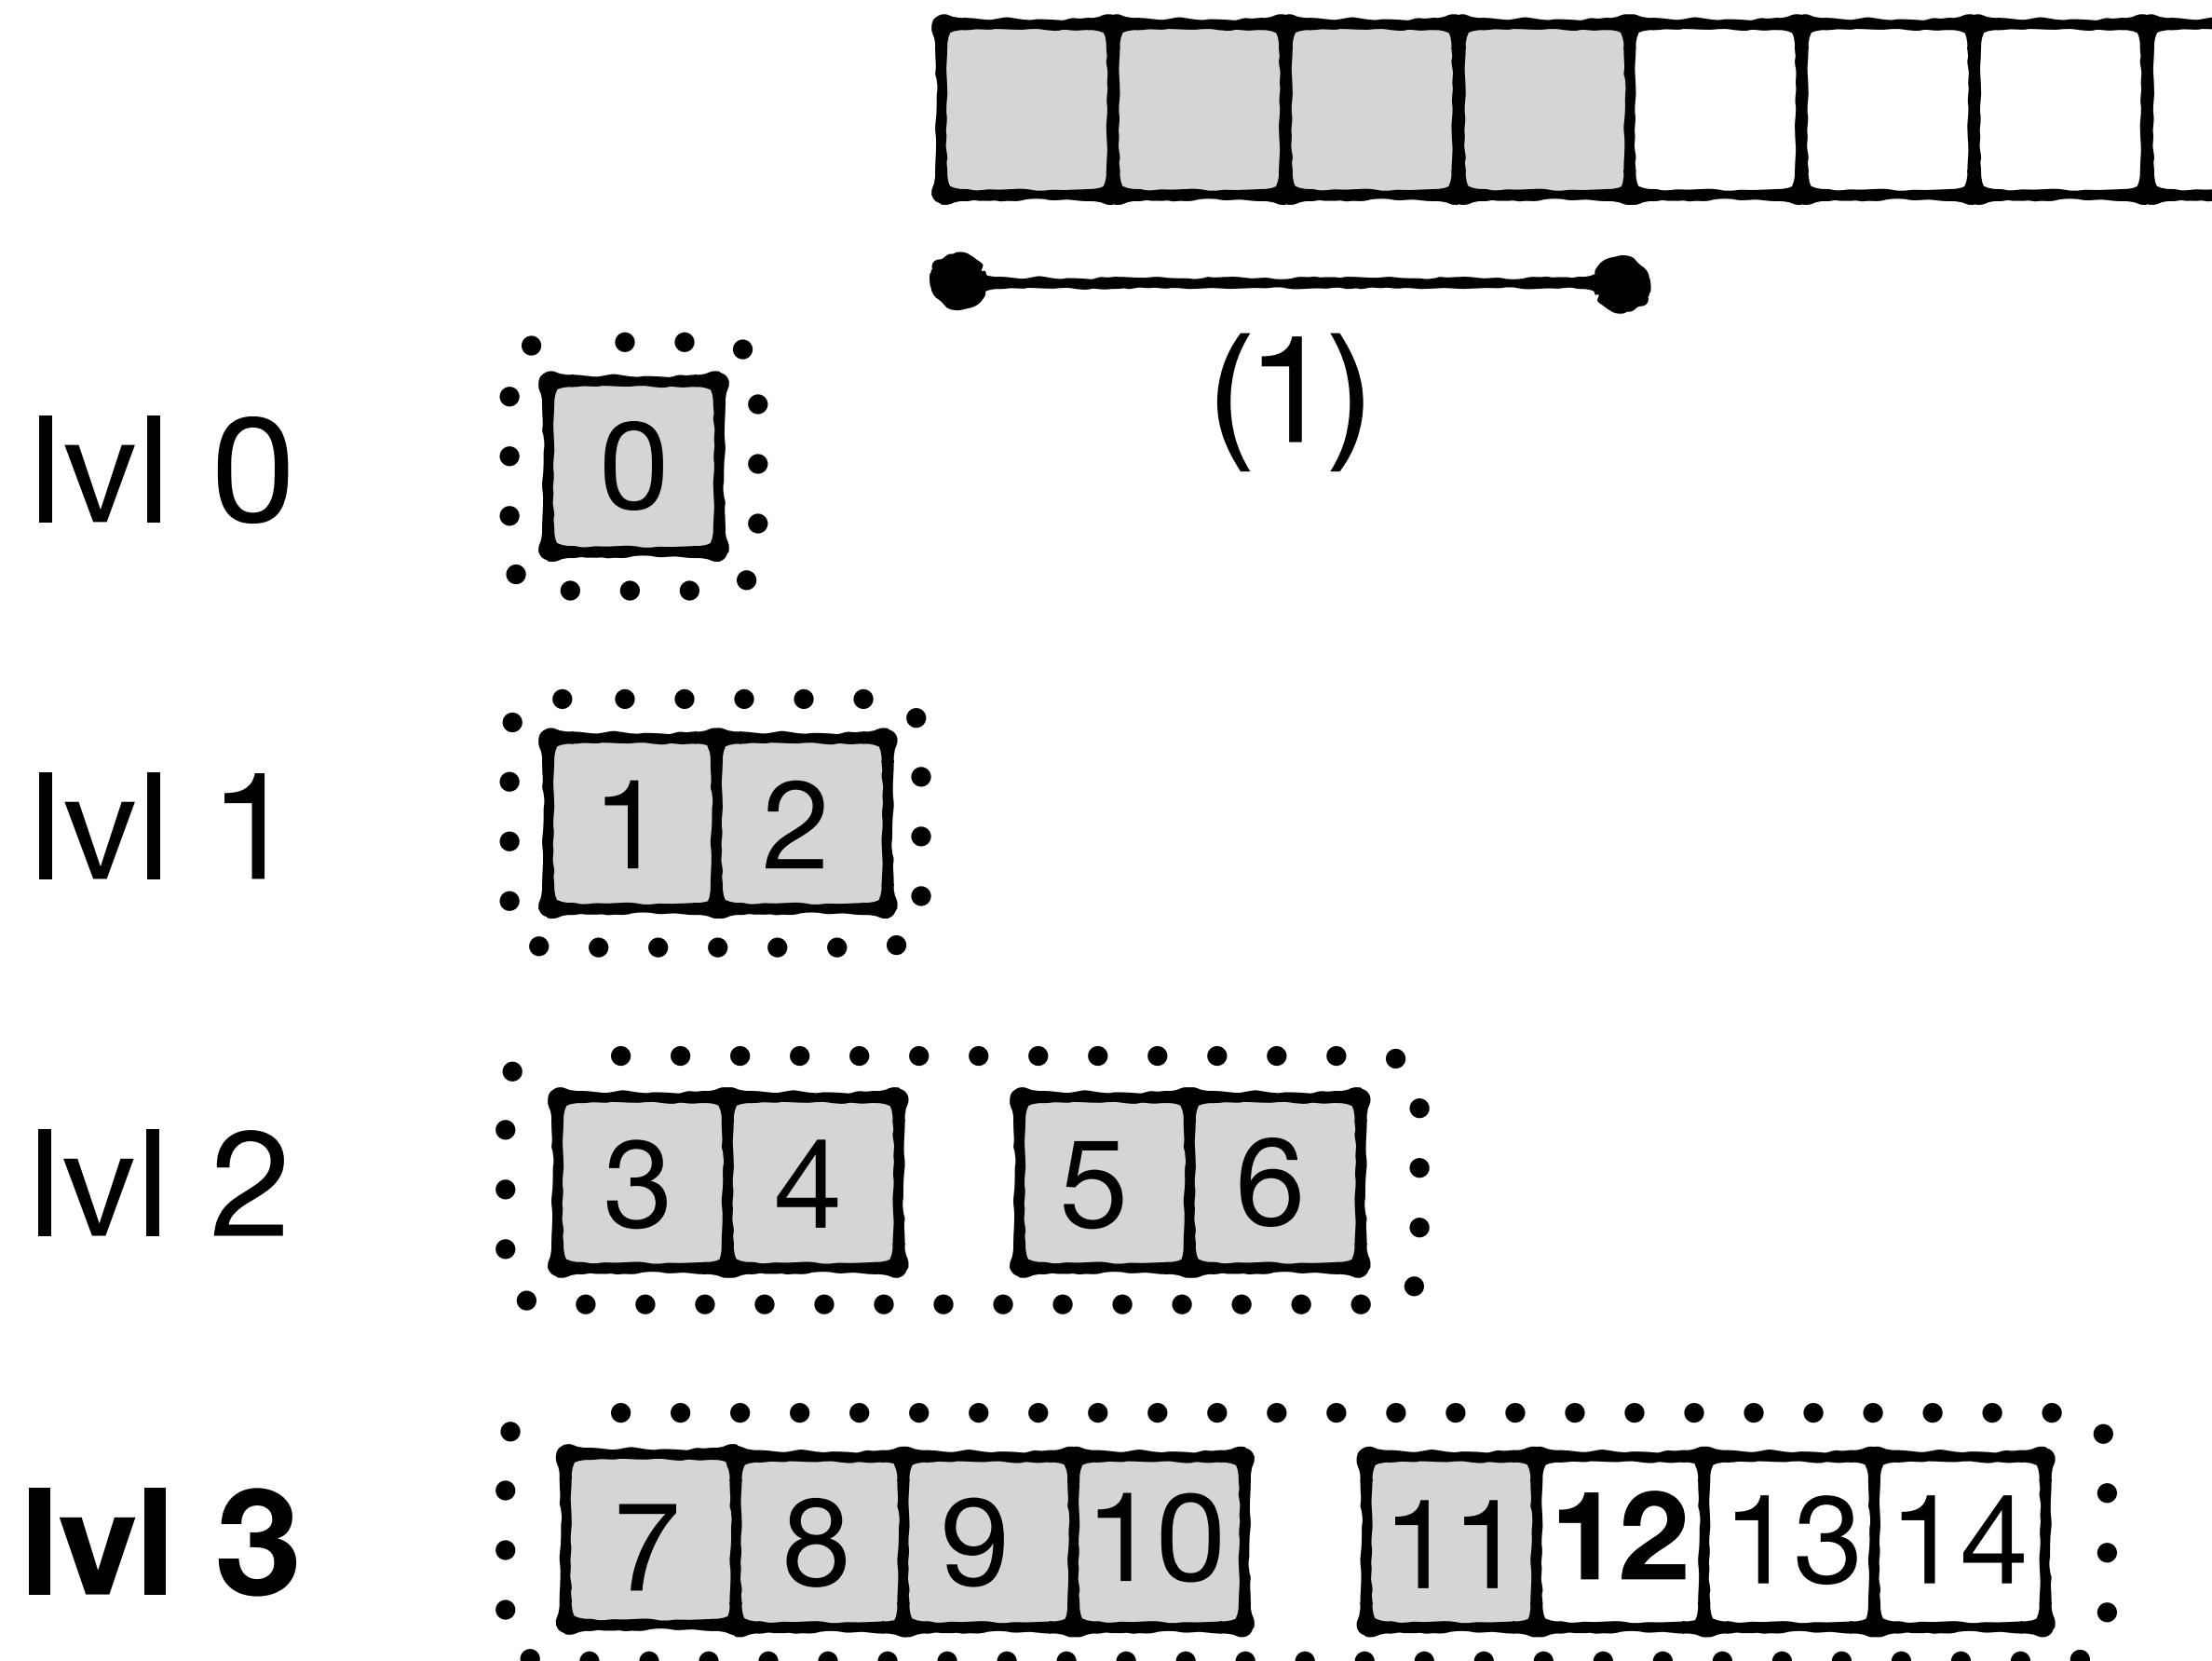
$$k = \lfloor \log_2(i + 1) \rfloor$$

$$= \text{sizeof}(i) - 1 - \text{clz}(i+1)$$



**Integer
length in bits**

get(i)



Identify target level k

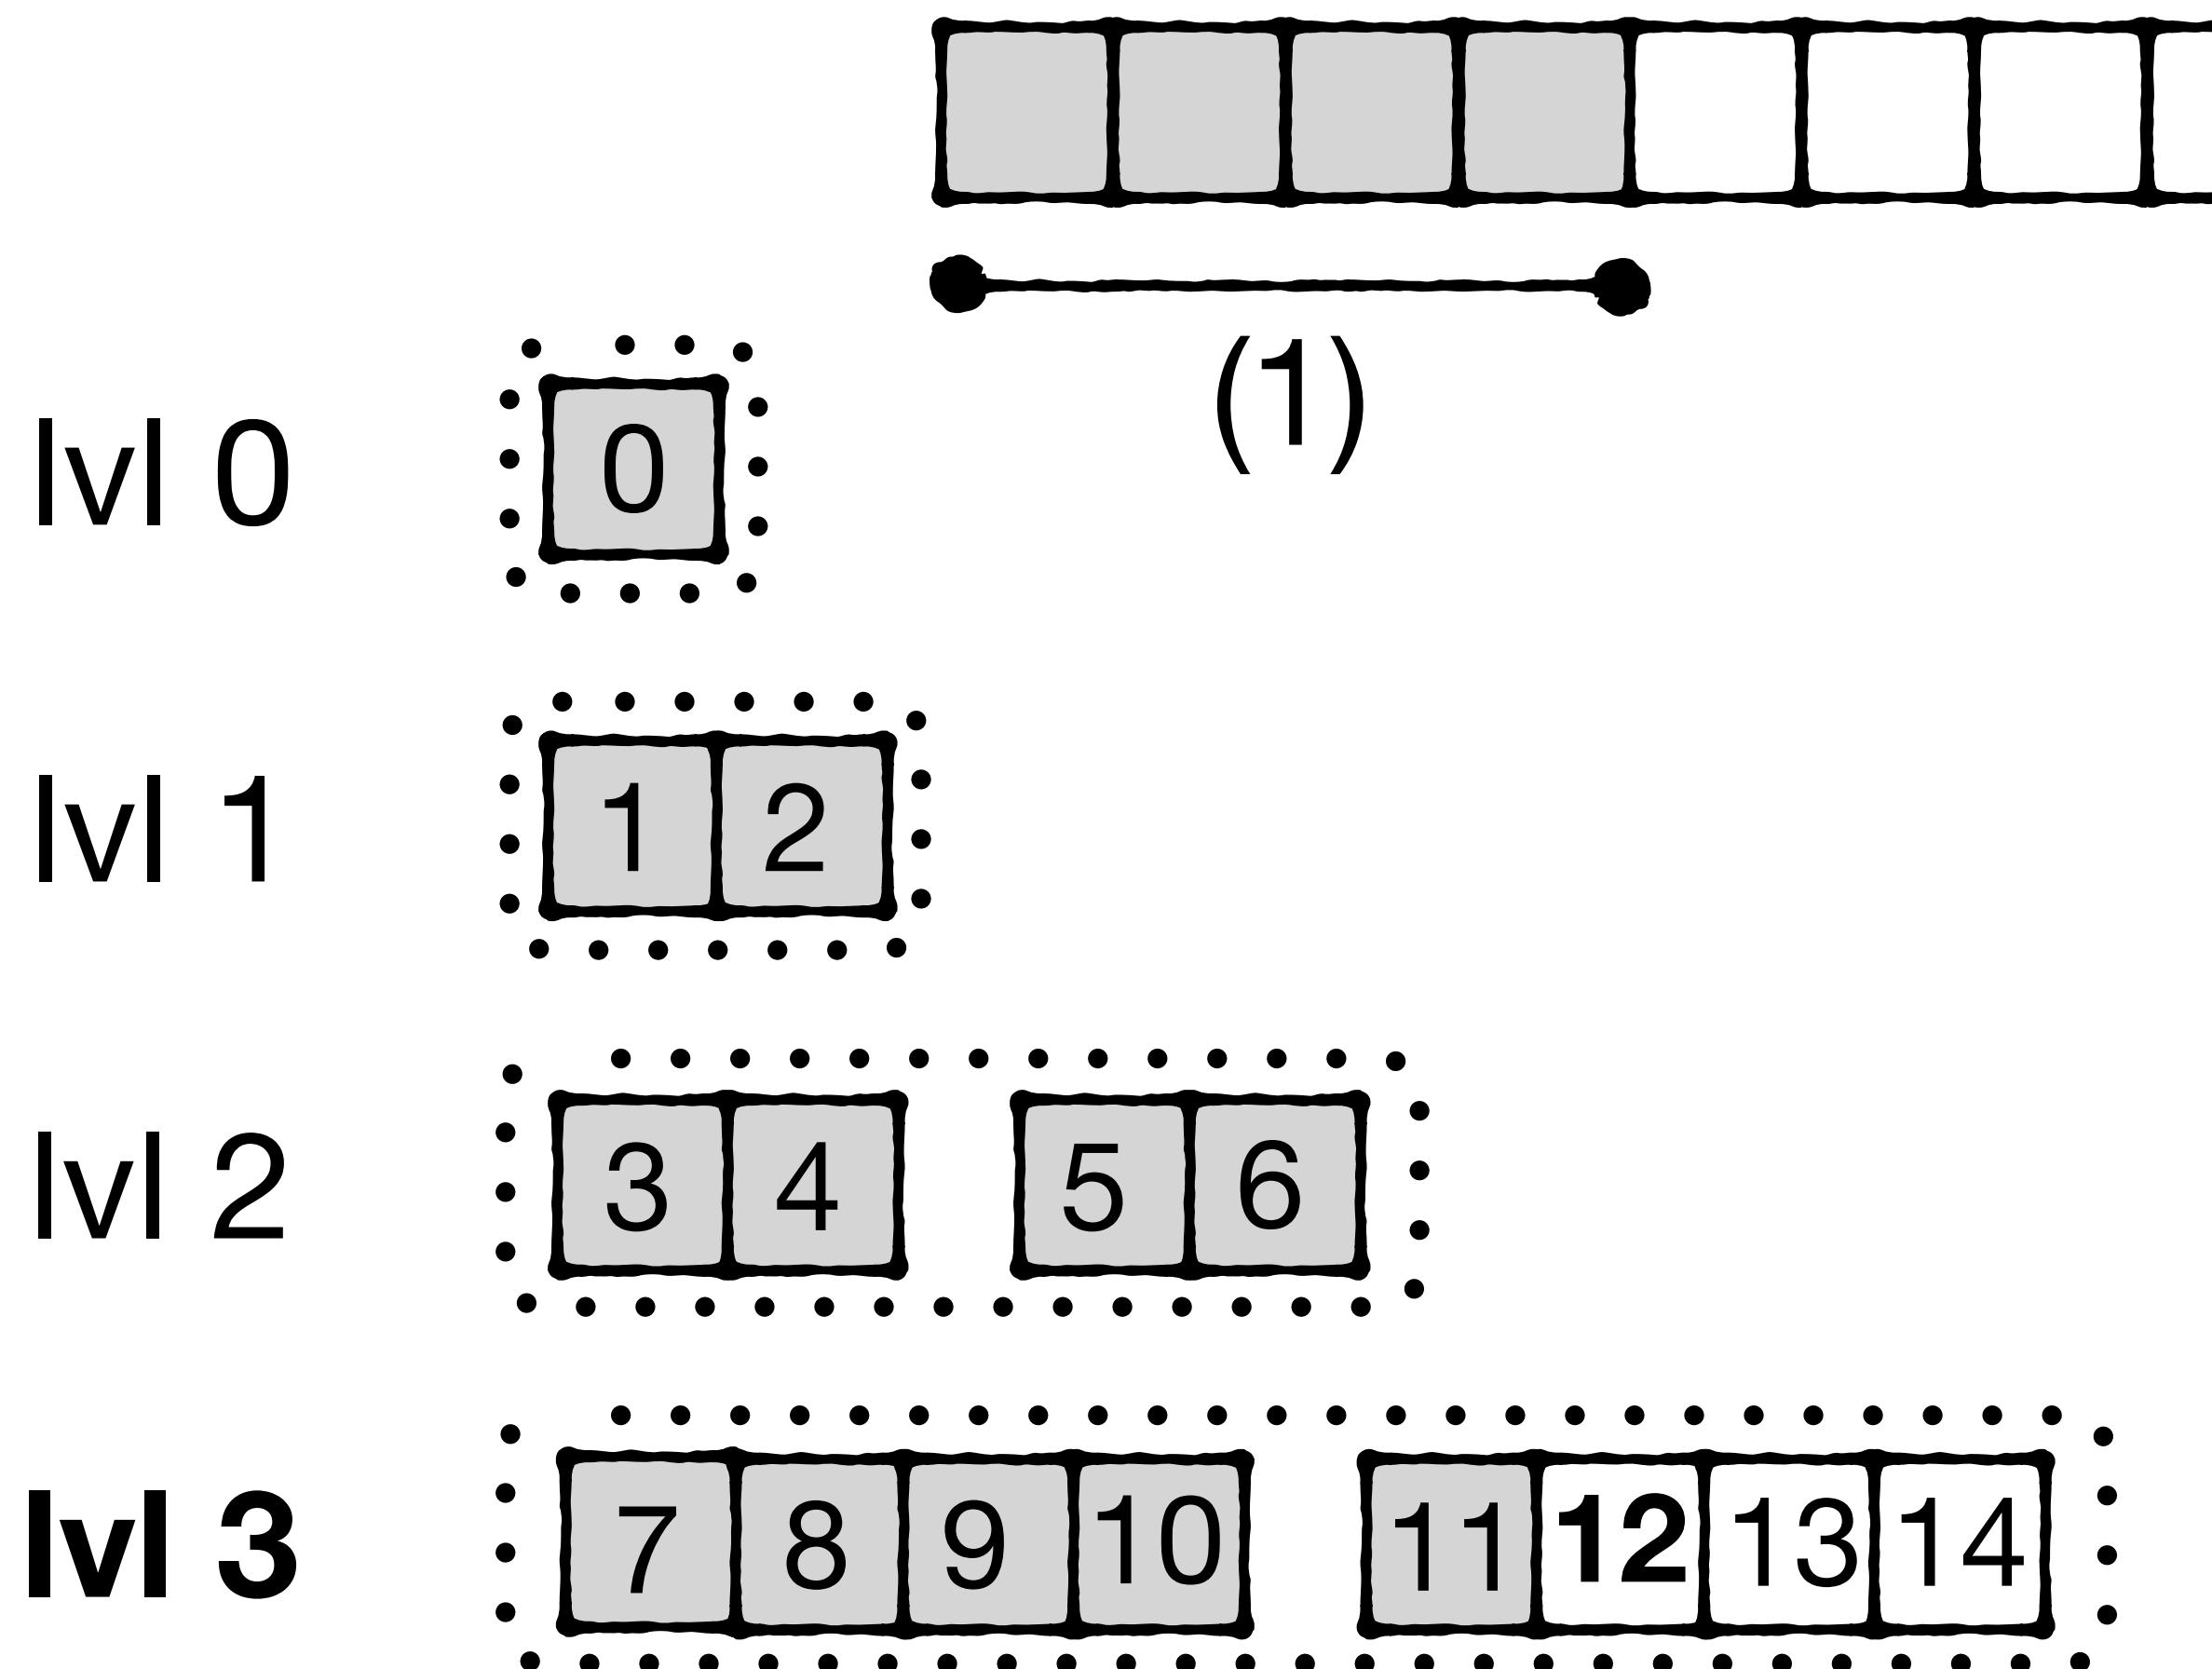
$$k = \lfloor \log_2(i + 1) \rfloor$$

$$= \text{sizeof}(i) - 1 - \text{clz}(i+1)$$



**Specialized CPU
command for #
leading zeros**

get(00001100)



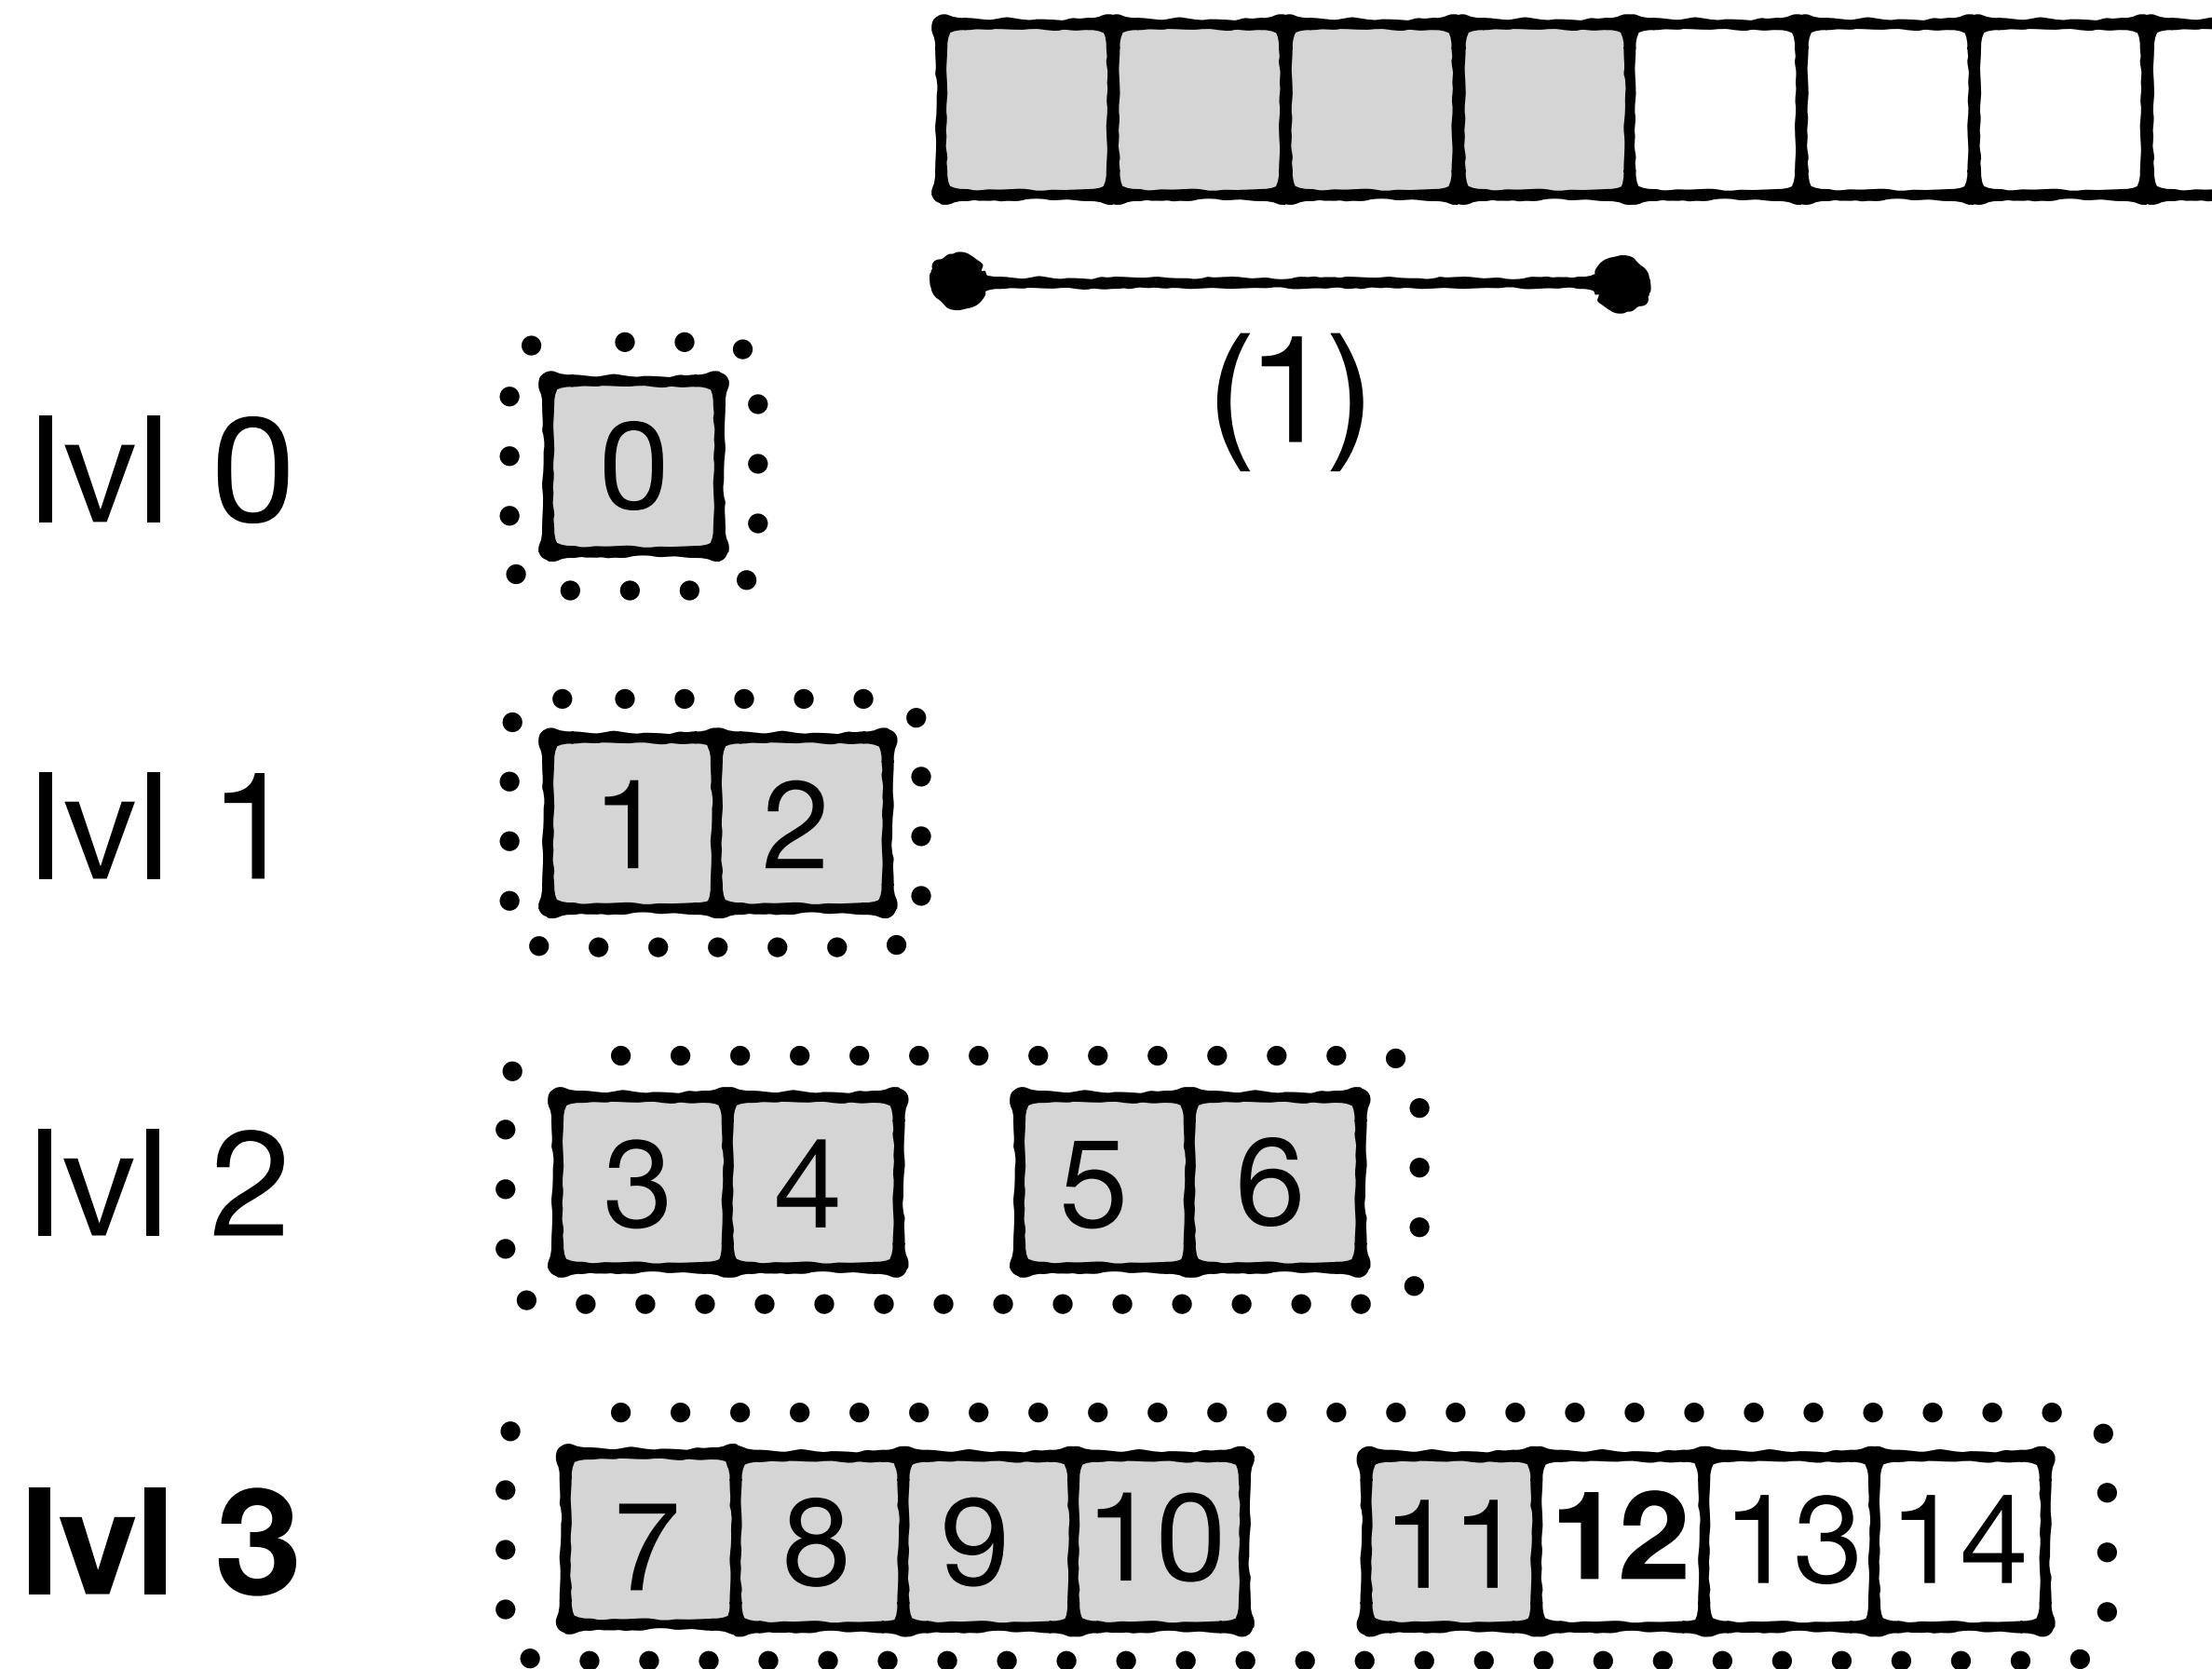
Identify target level k

$$k = \lfloor \log_2(i + 1) \rfloor$$

$$= \text{sizeof}(i) - 1 - \text{clz}(i+1)$$

$$8 - 1 - 4 = 3$$

get(i)



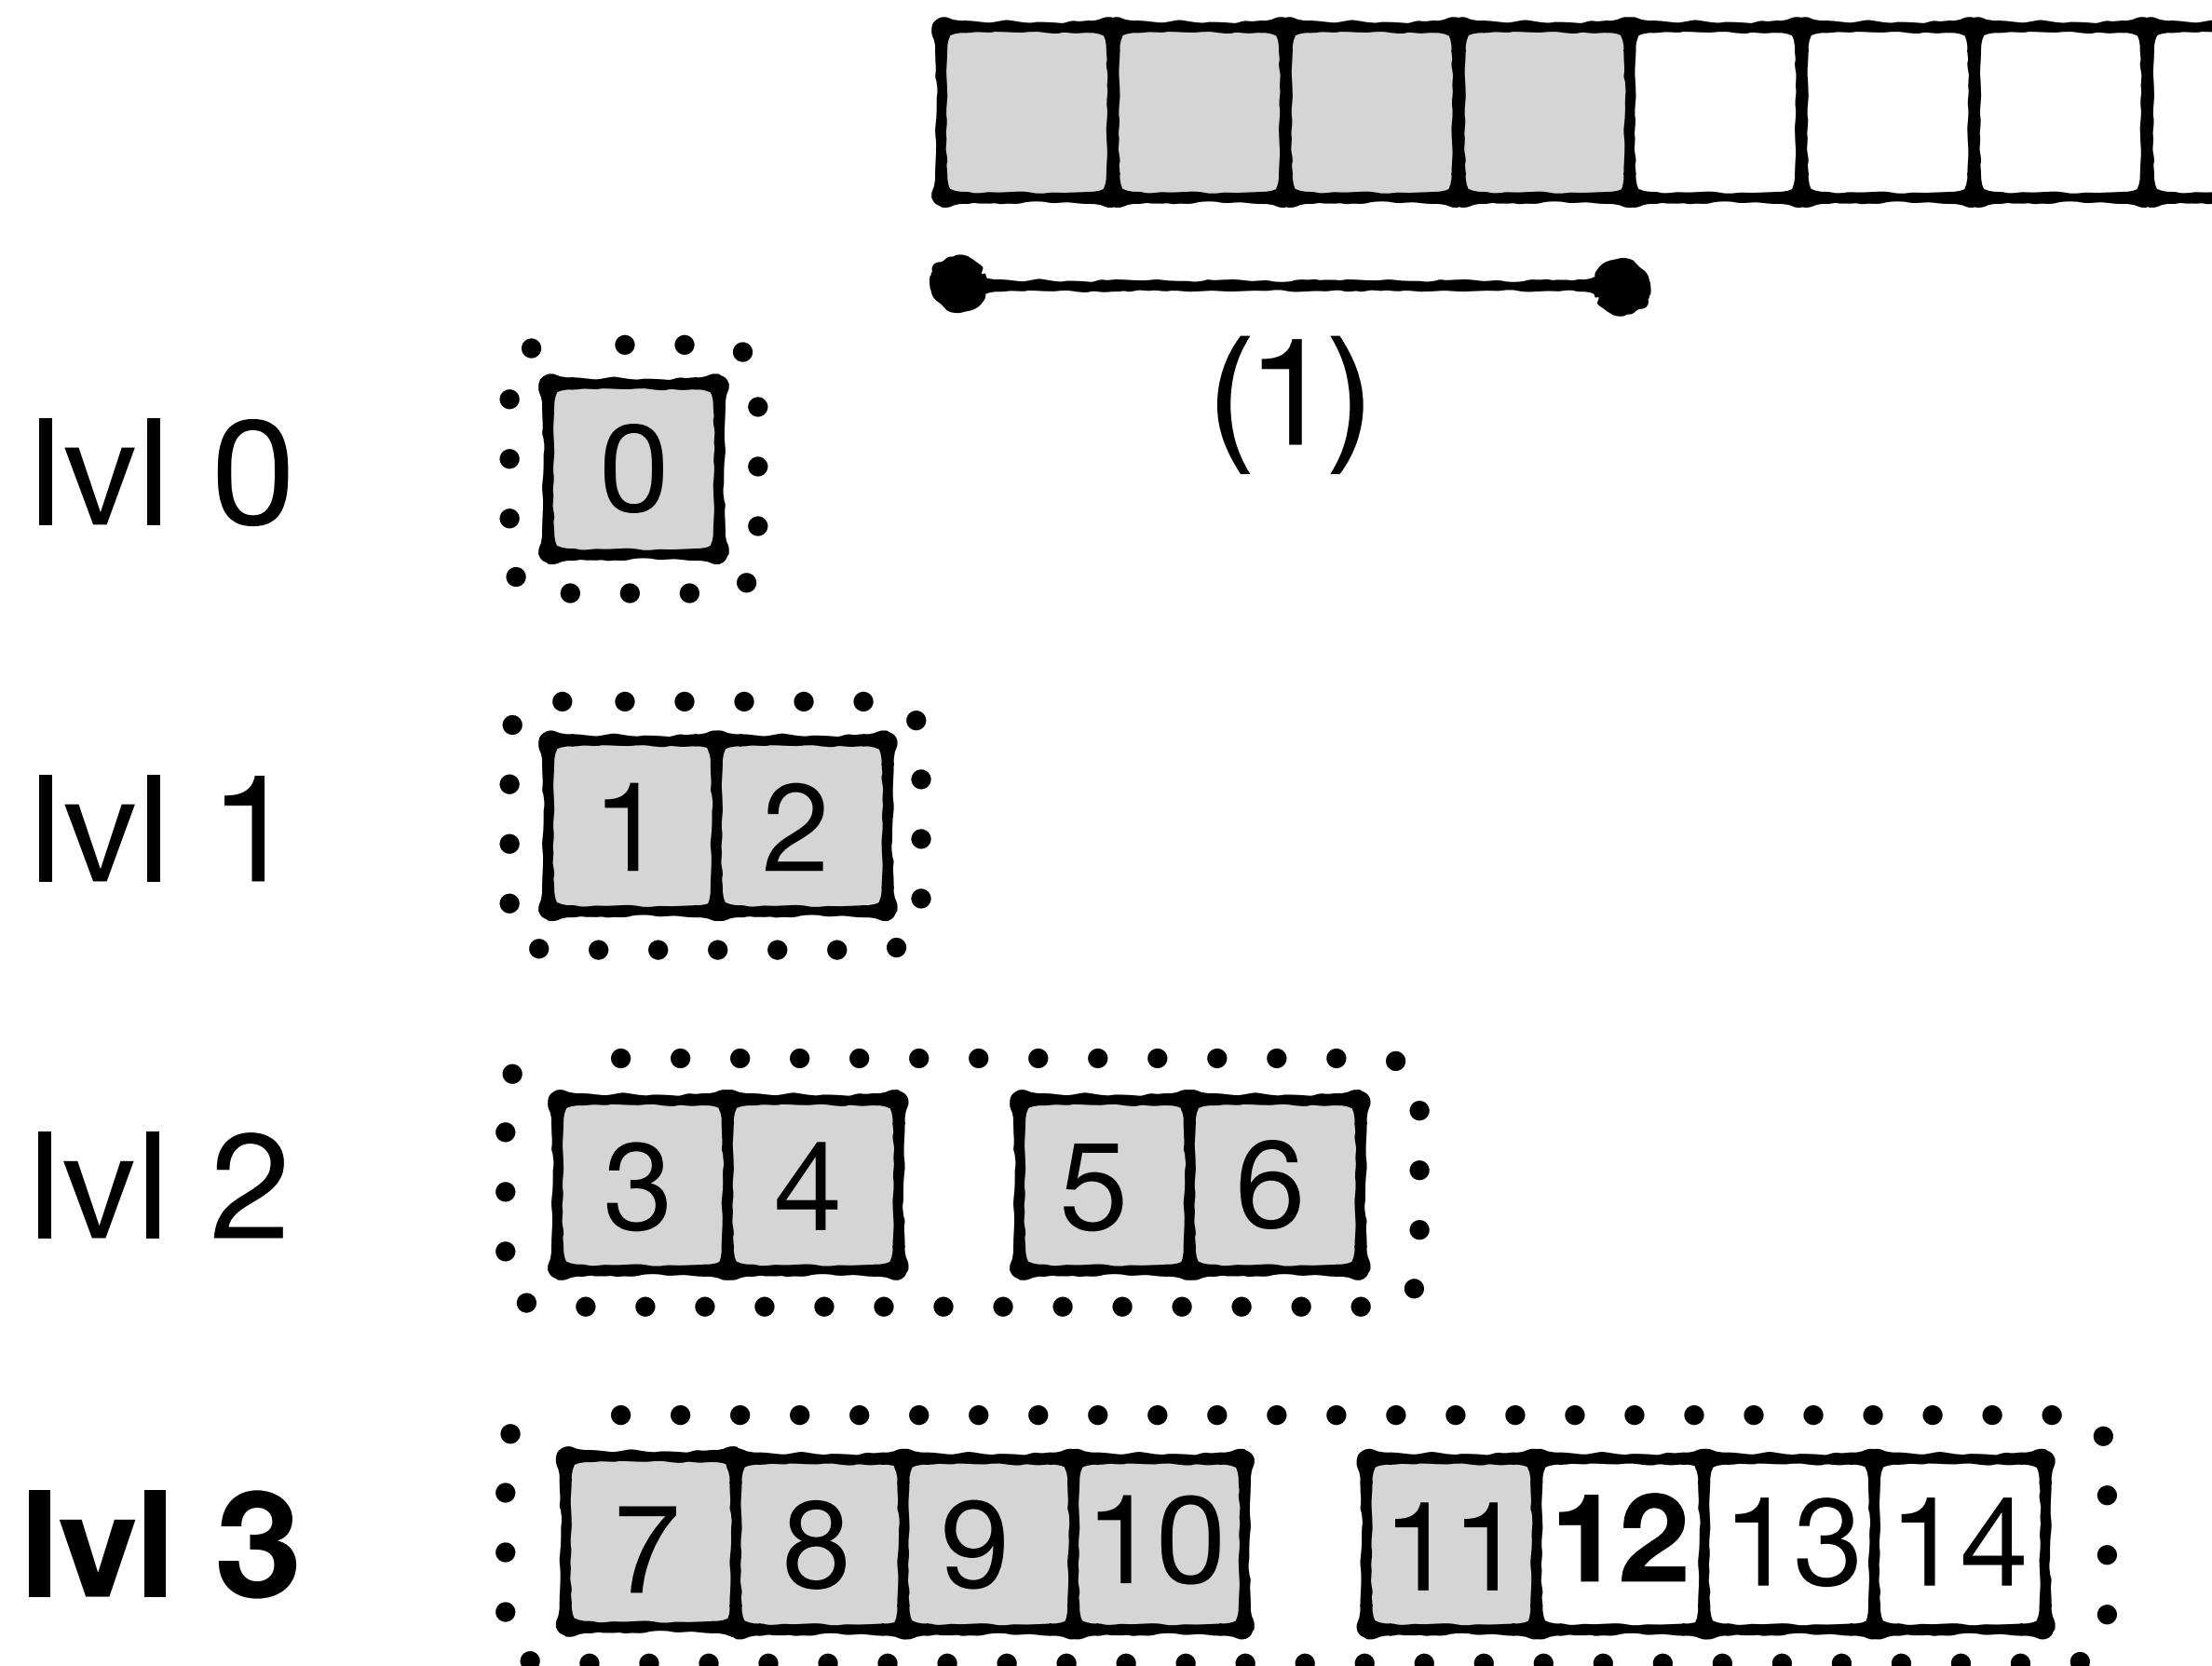
Identify target level k

$$k = \lfloor \log_2(i + 1) \rfloor$$

$$= \text{sizeof}(i) - 1 - \text{clz}(i+1)$$

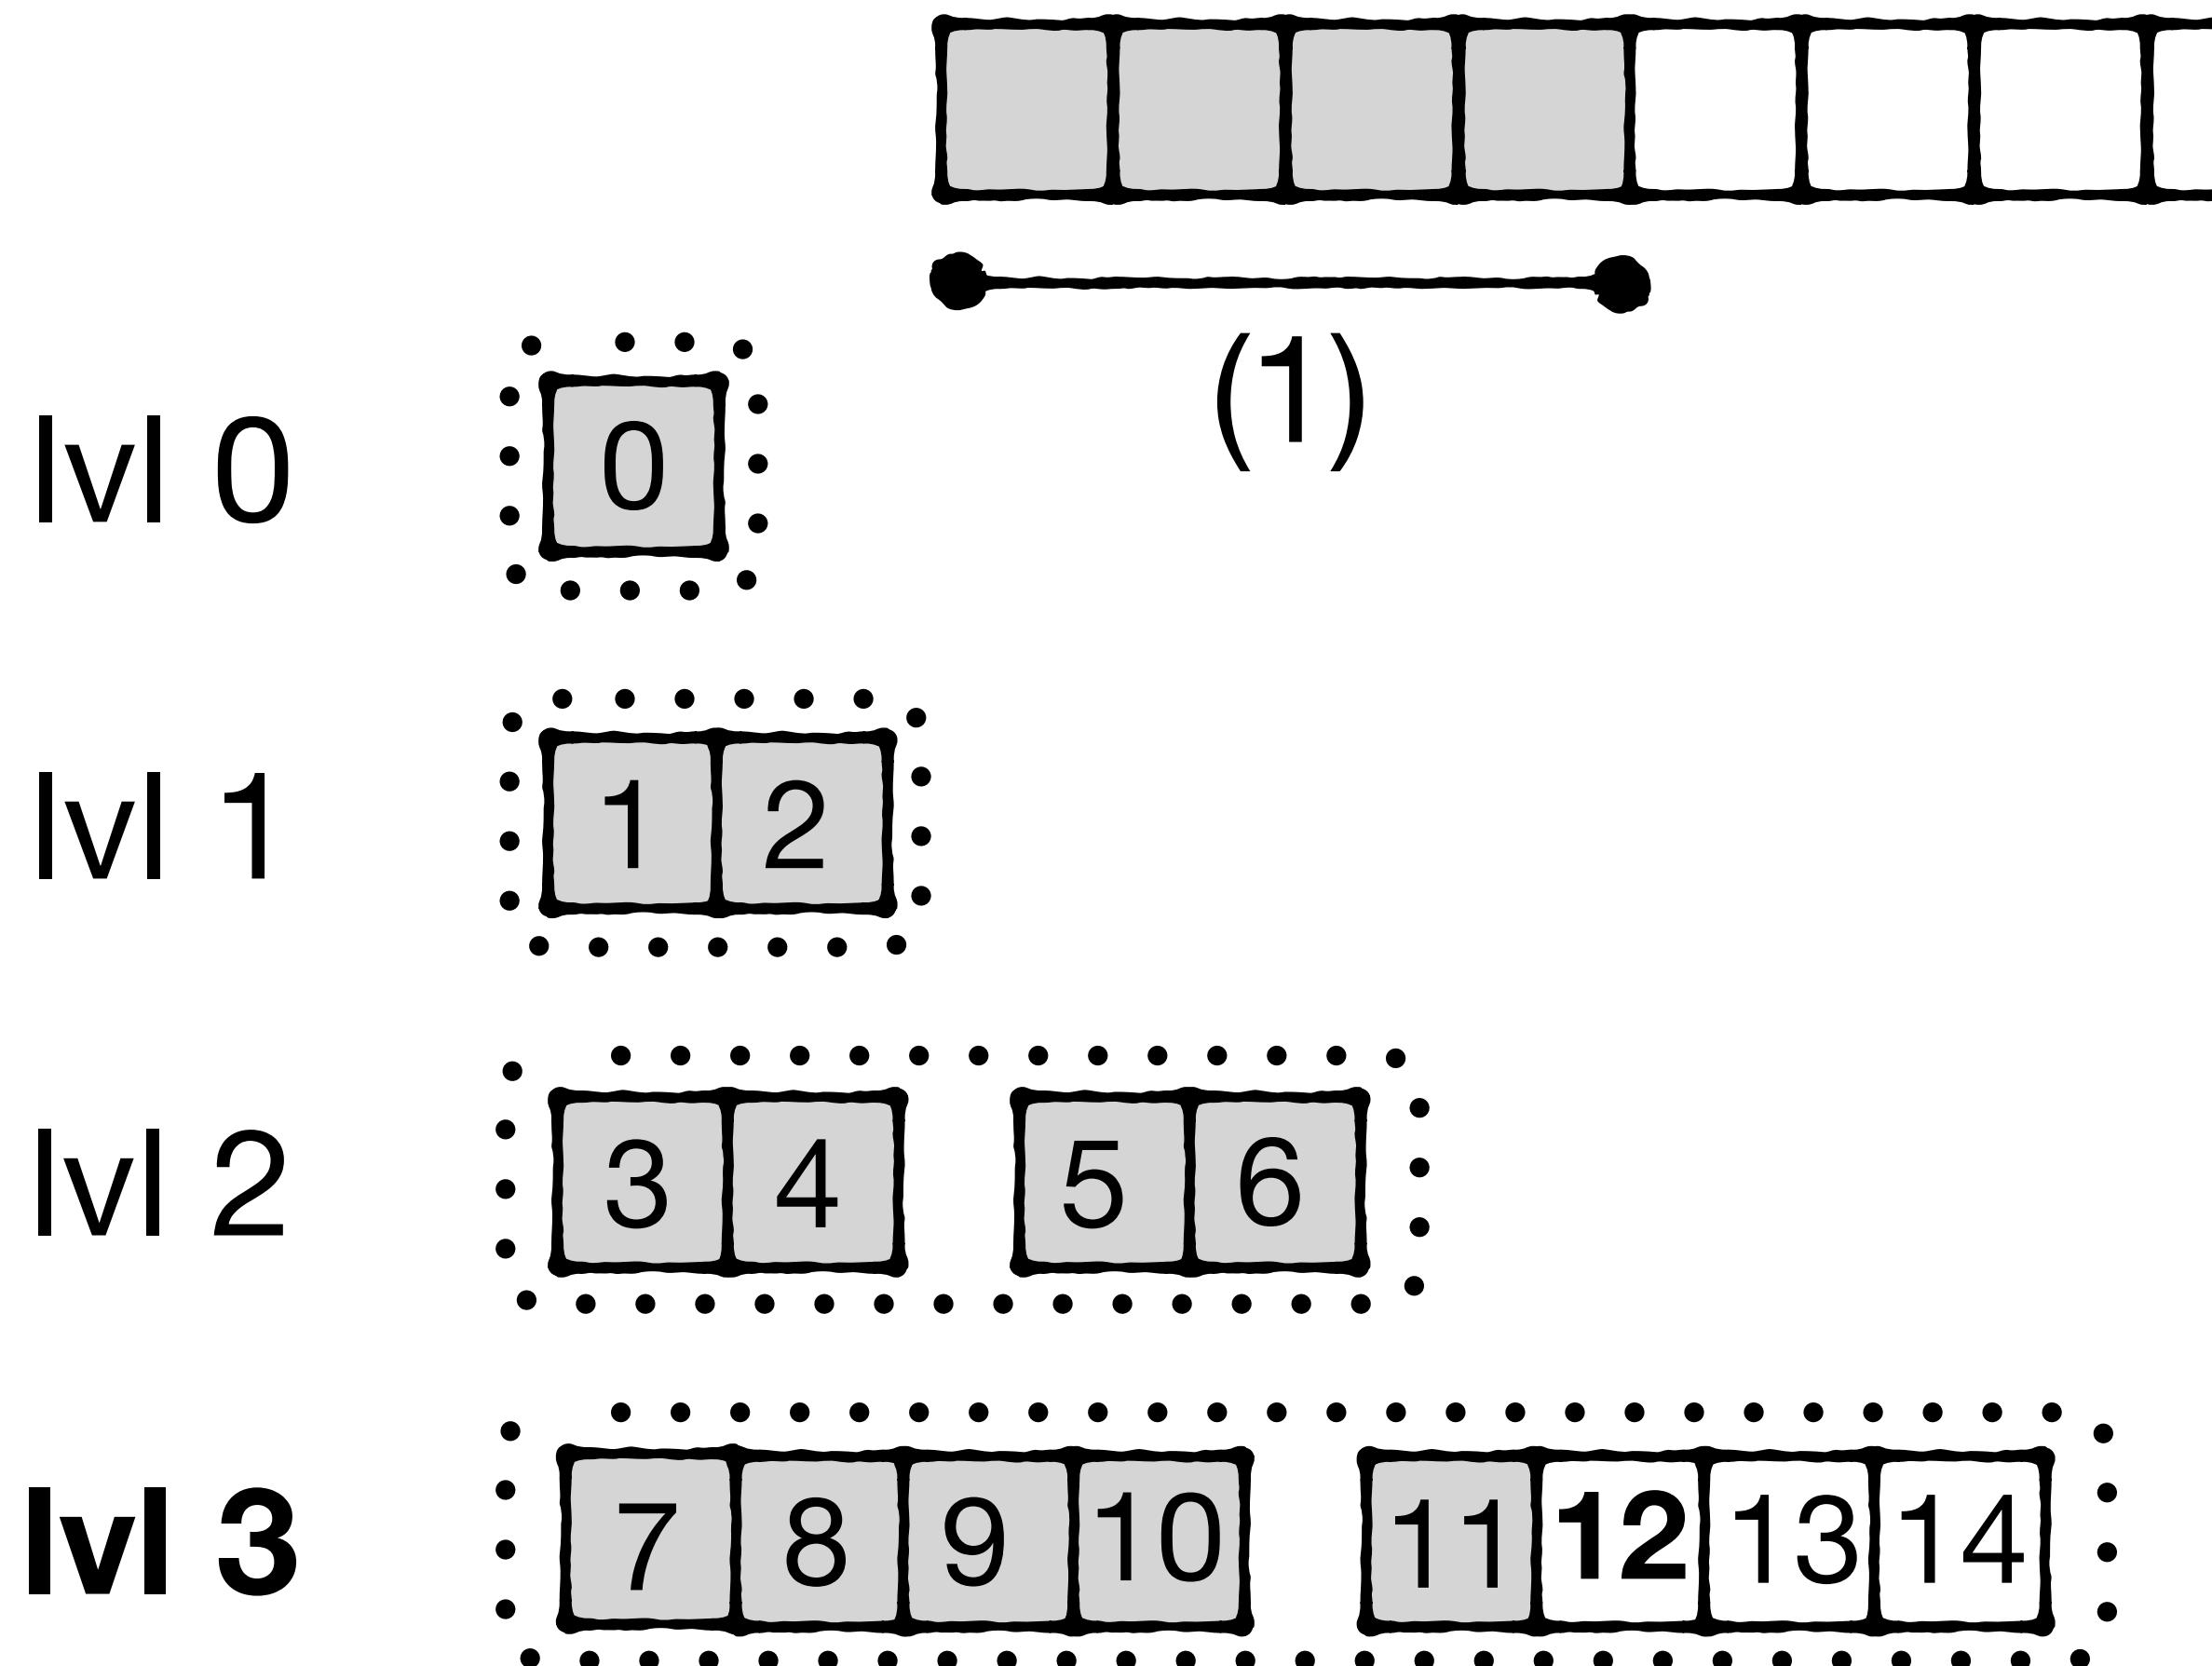
≈1 ns rather than ≈7ns

get(i)



blocks in levels 0 to k-1?

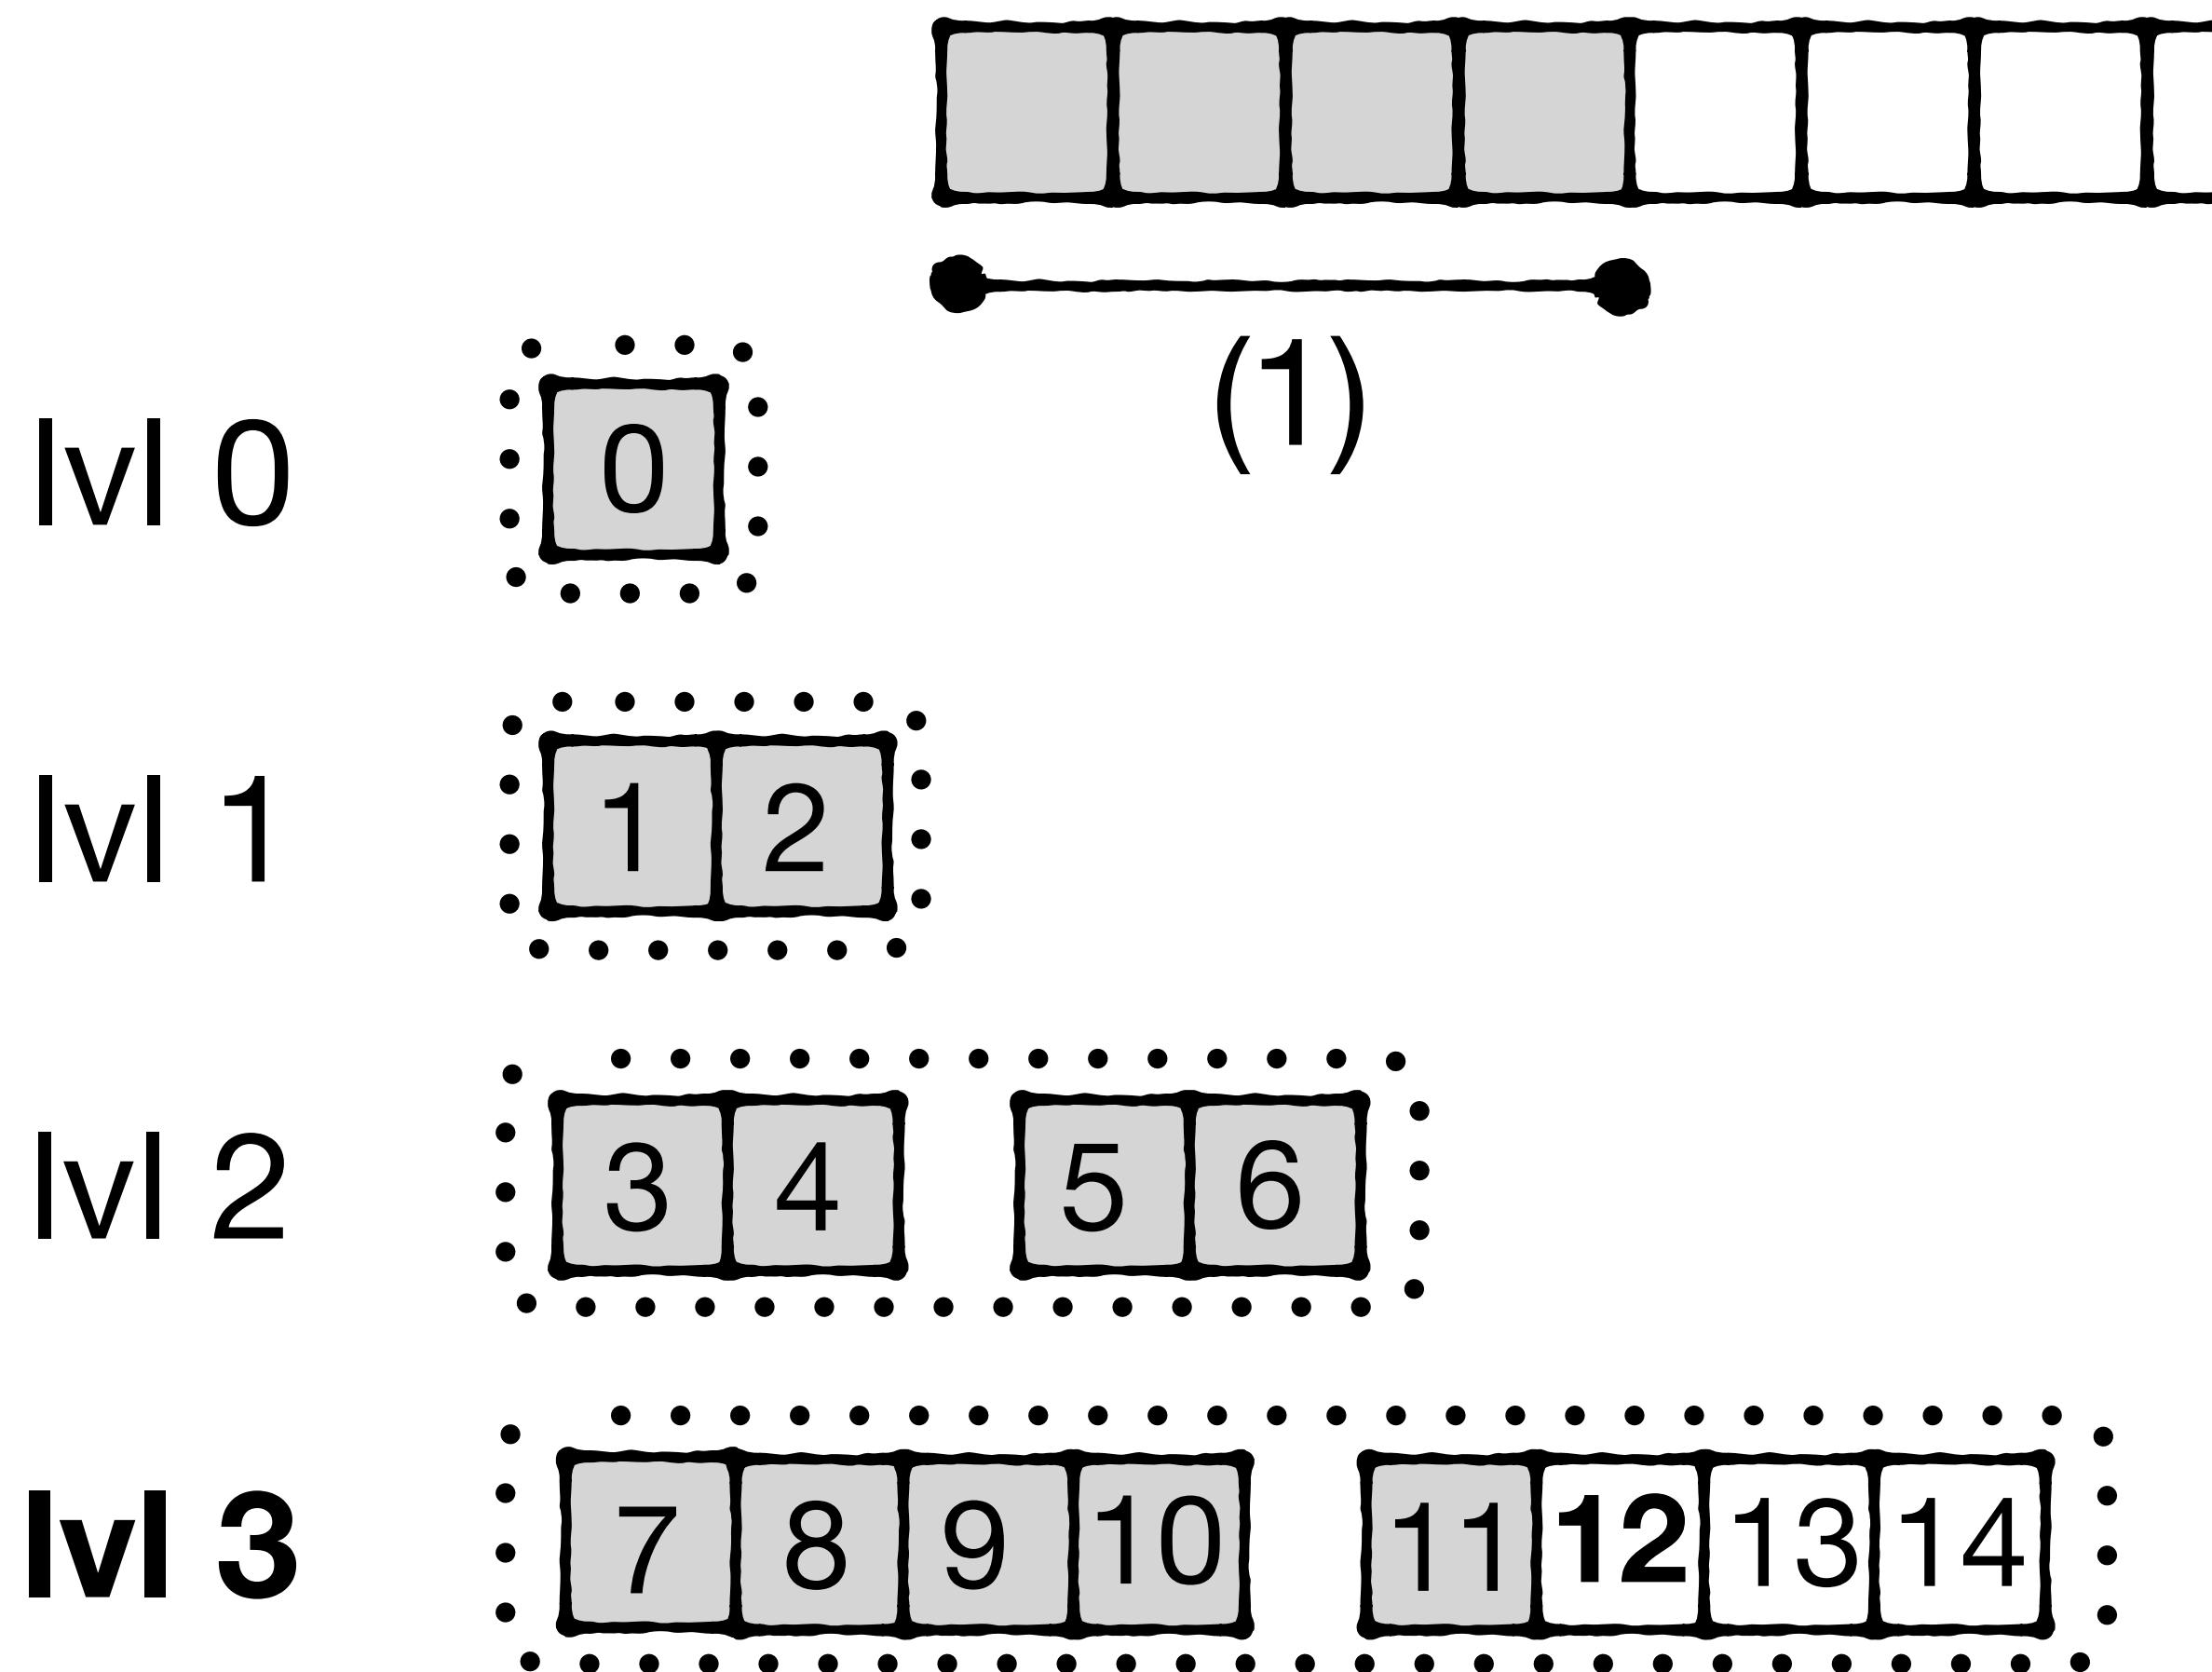
get(i)



blocks in levels 0 to k-1

$$= 2^{\lfloor k/2 \rfloor} \cdot (2 + (k \bmod 2)) - 2$$

get(i)

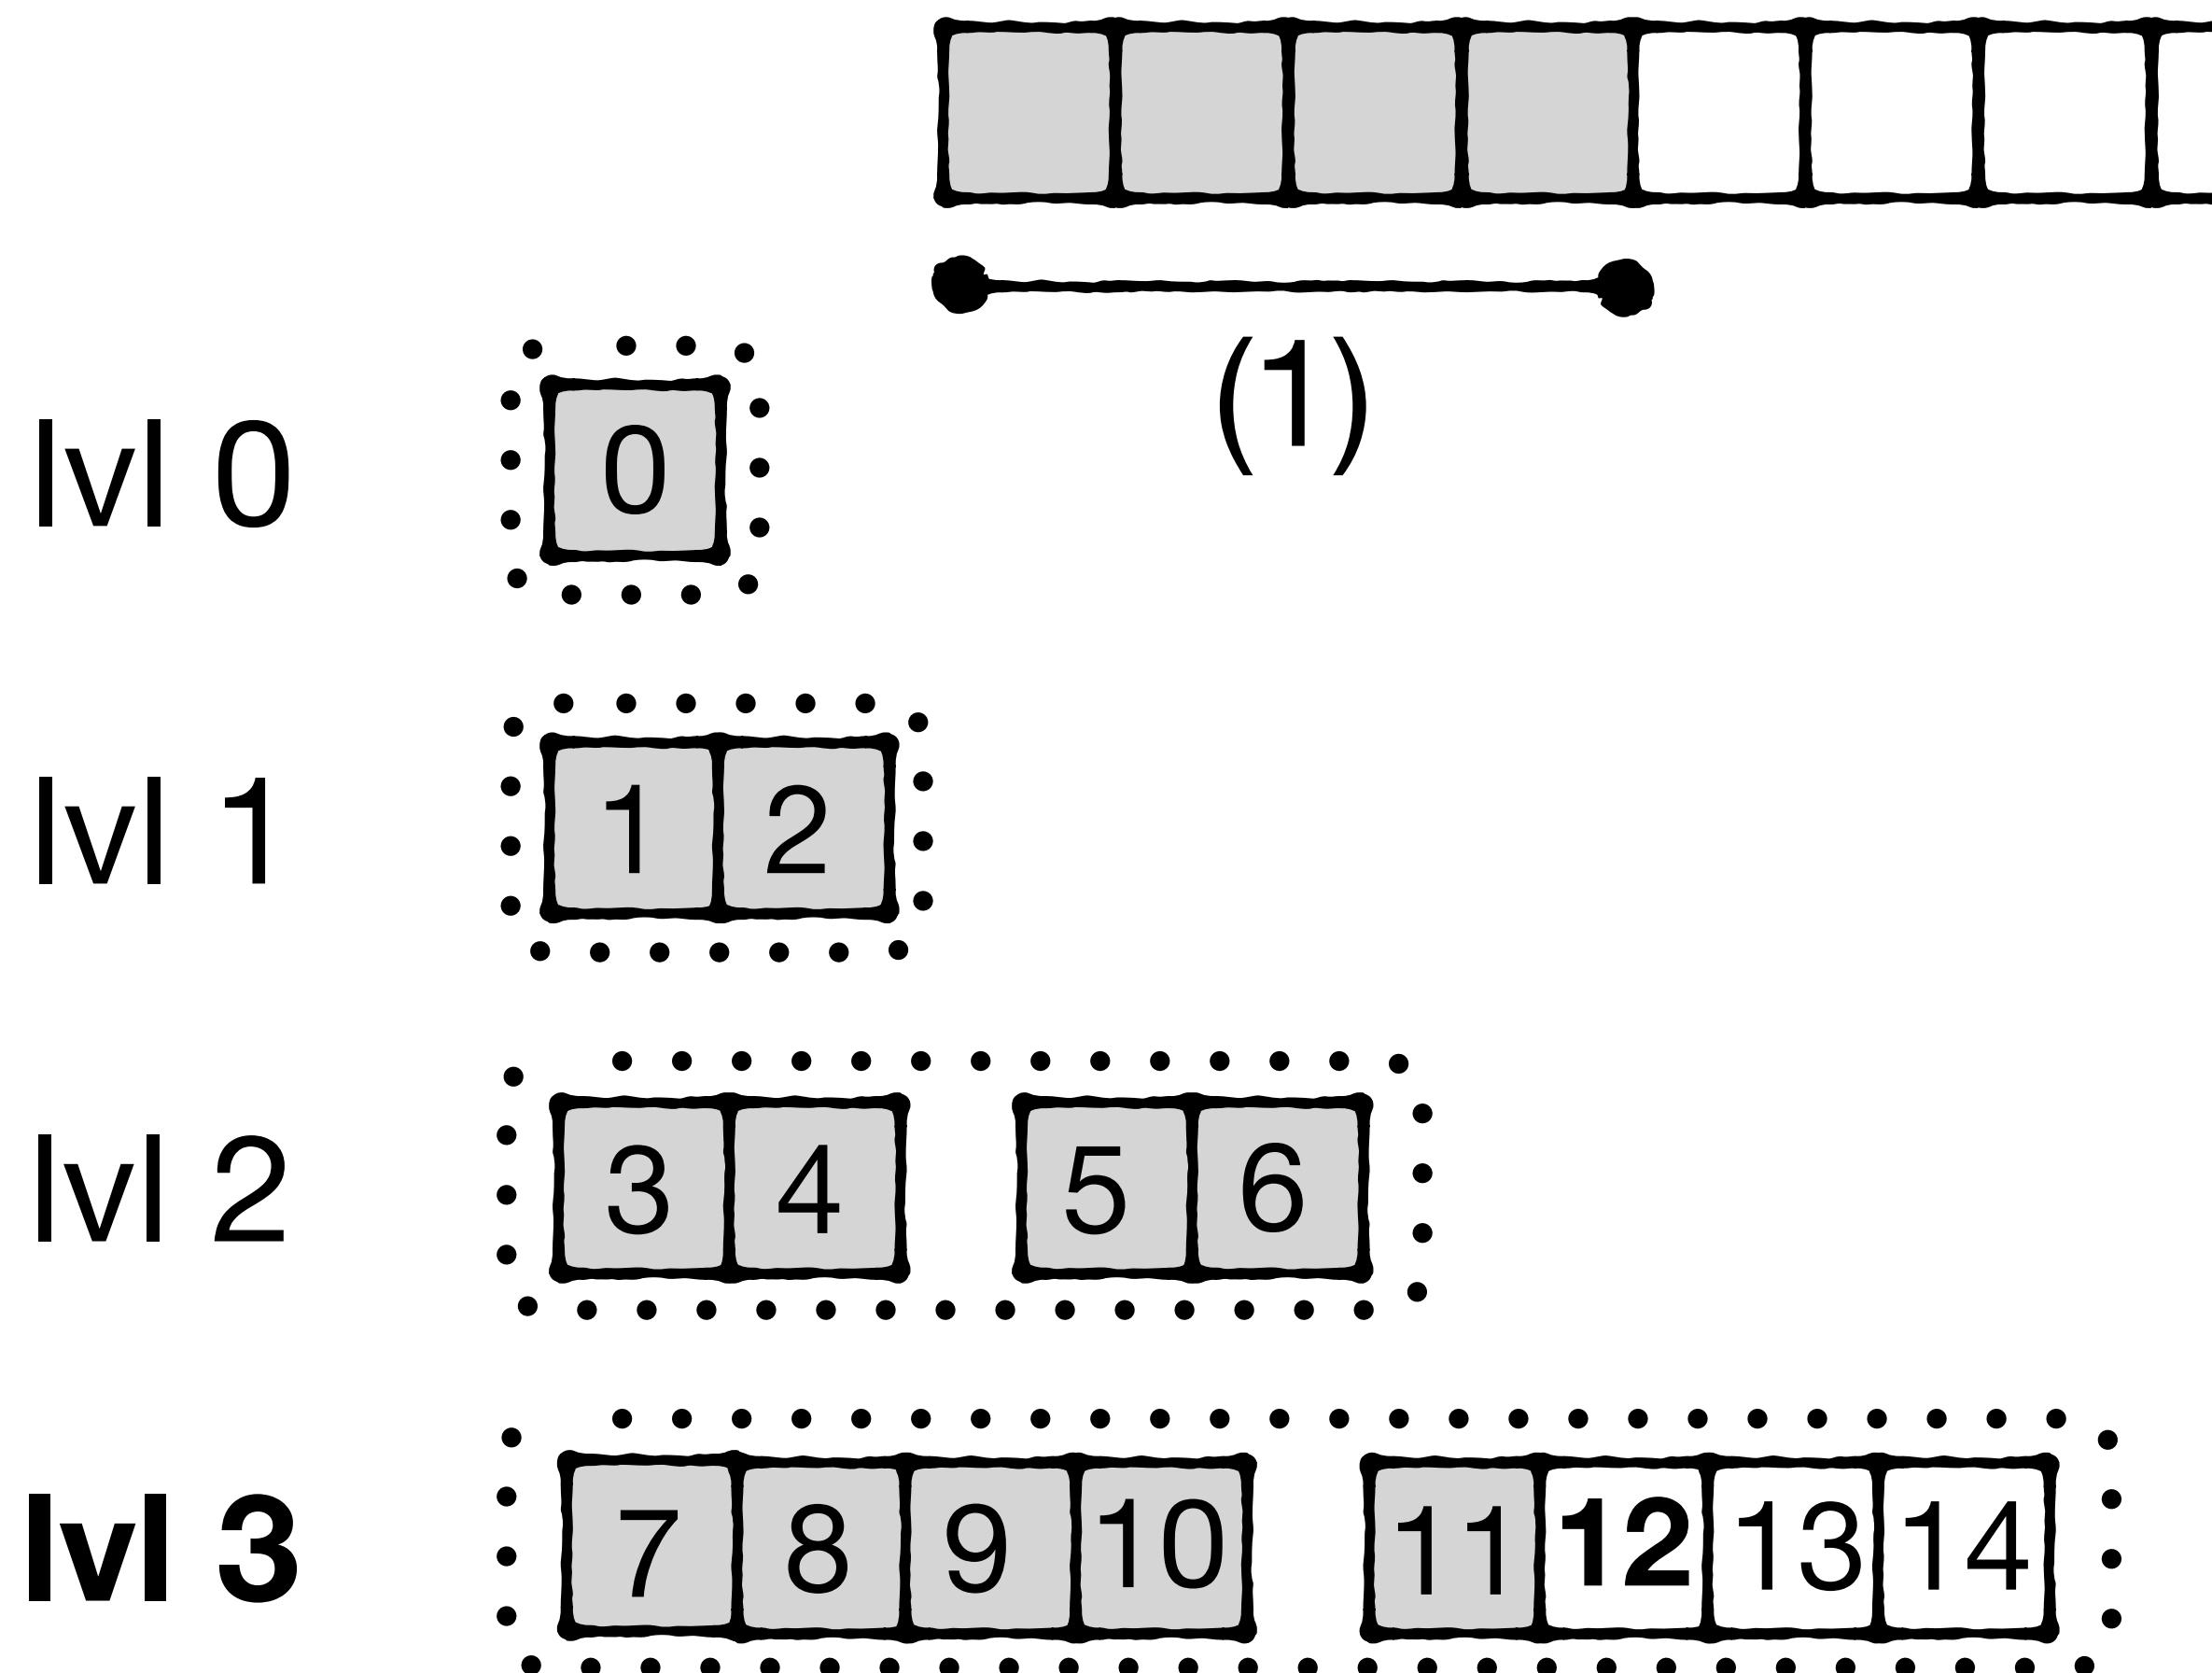


blocks in levels 0 to k-1

$$= 2^{\lfloor k/2 \rfloor} \cdot (2 + (k \bmod 2)) - 2$$

Original paper gets this wrong, fixed credit to Hyuhng Min

get(i)

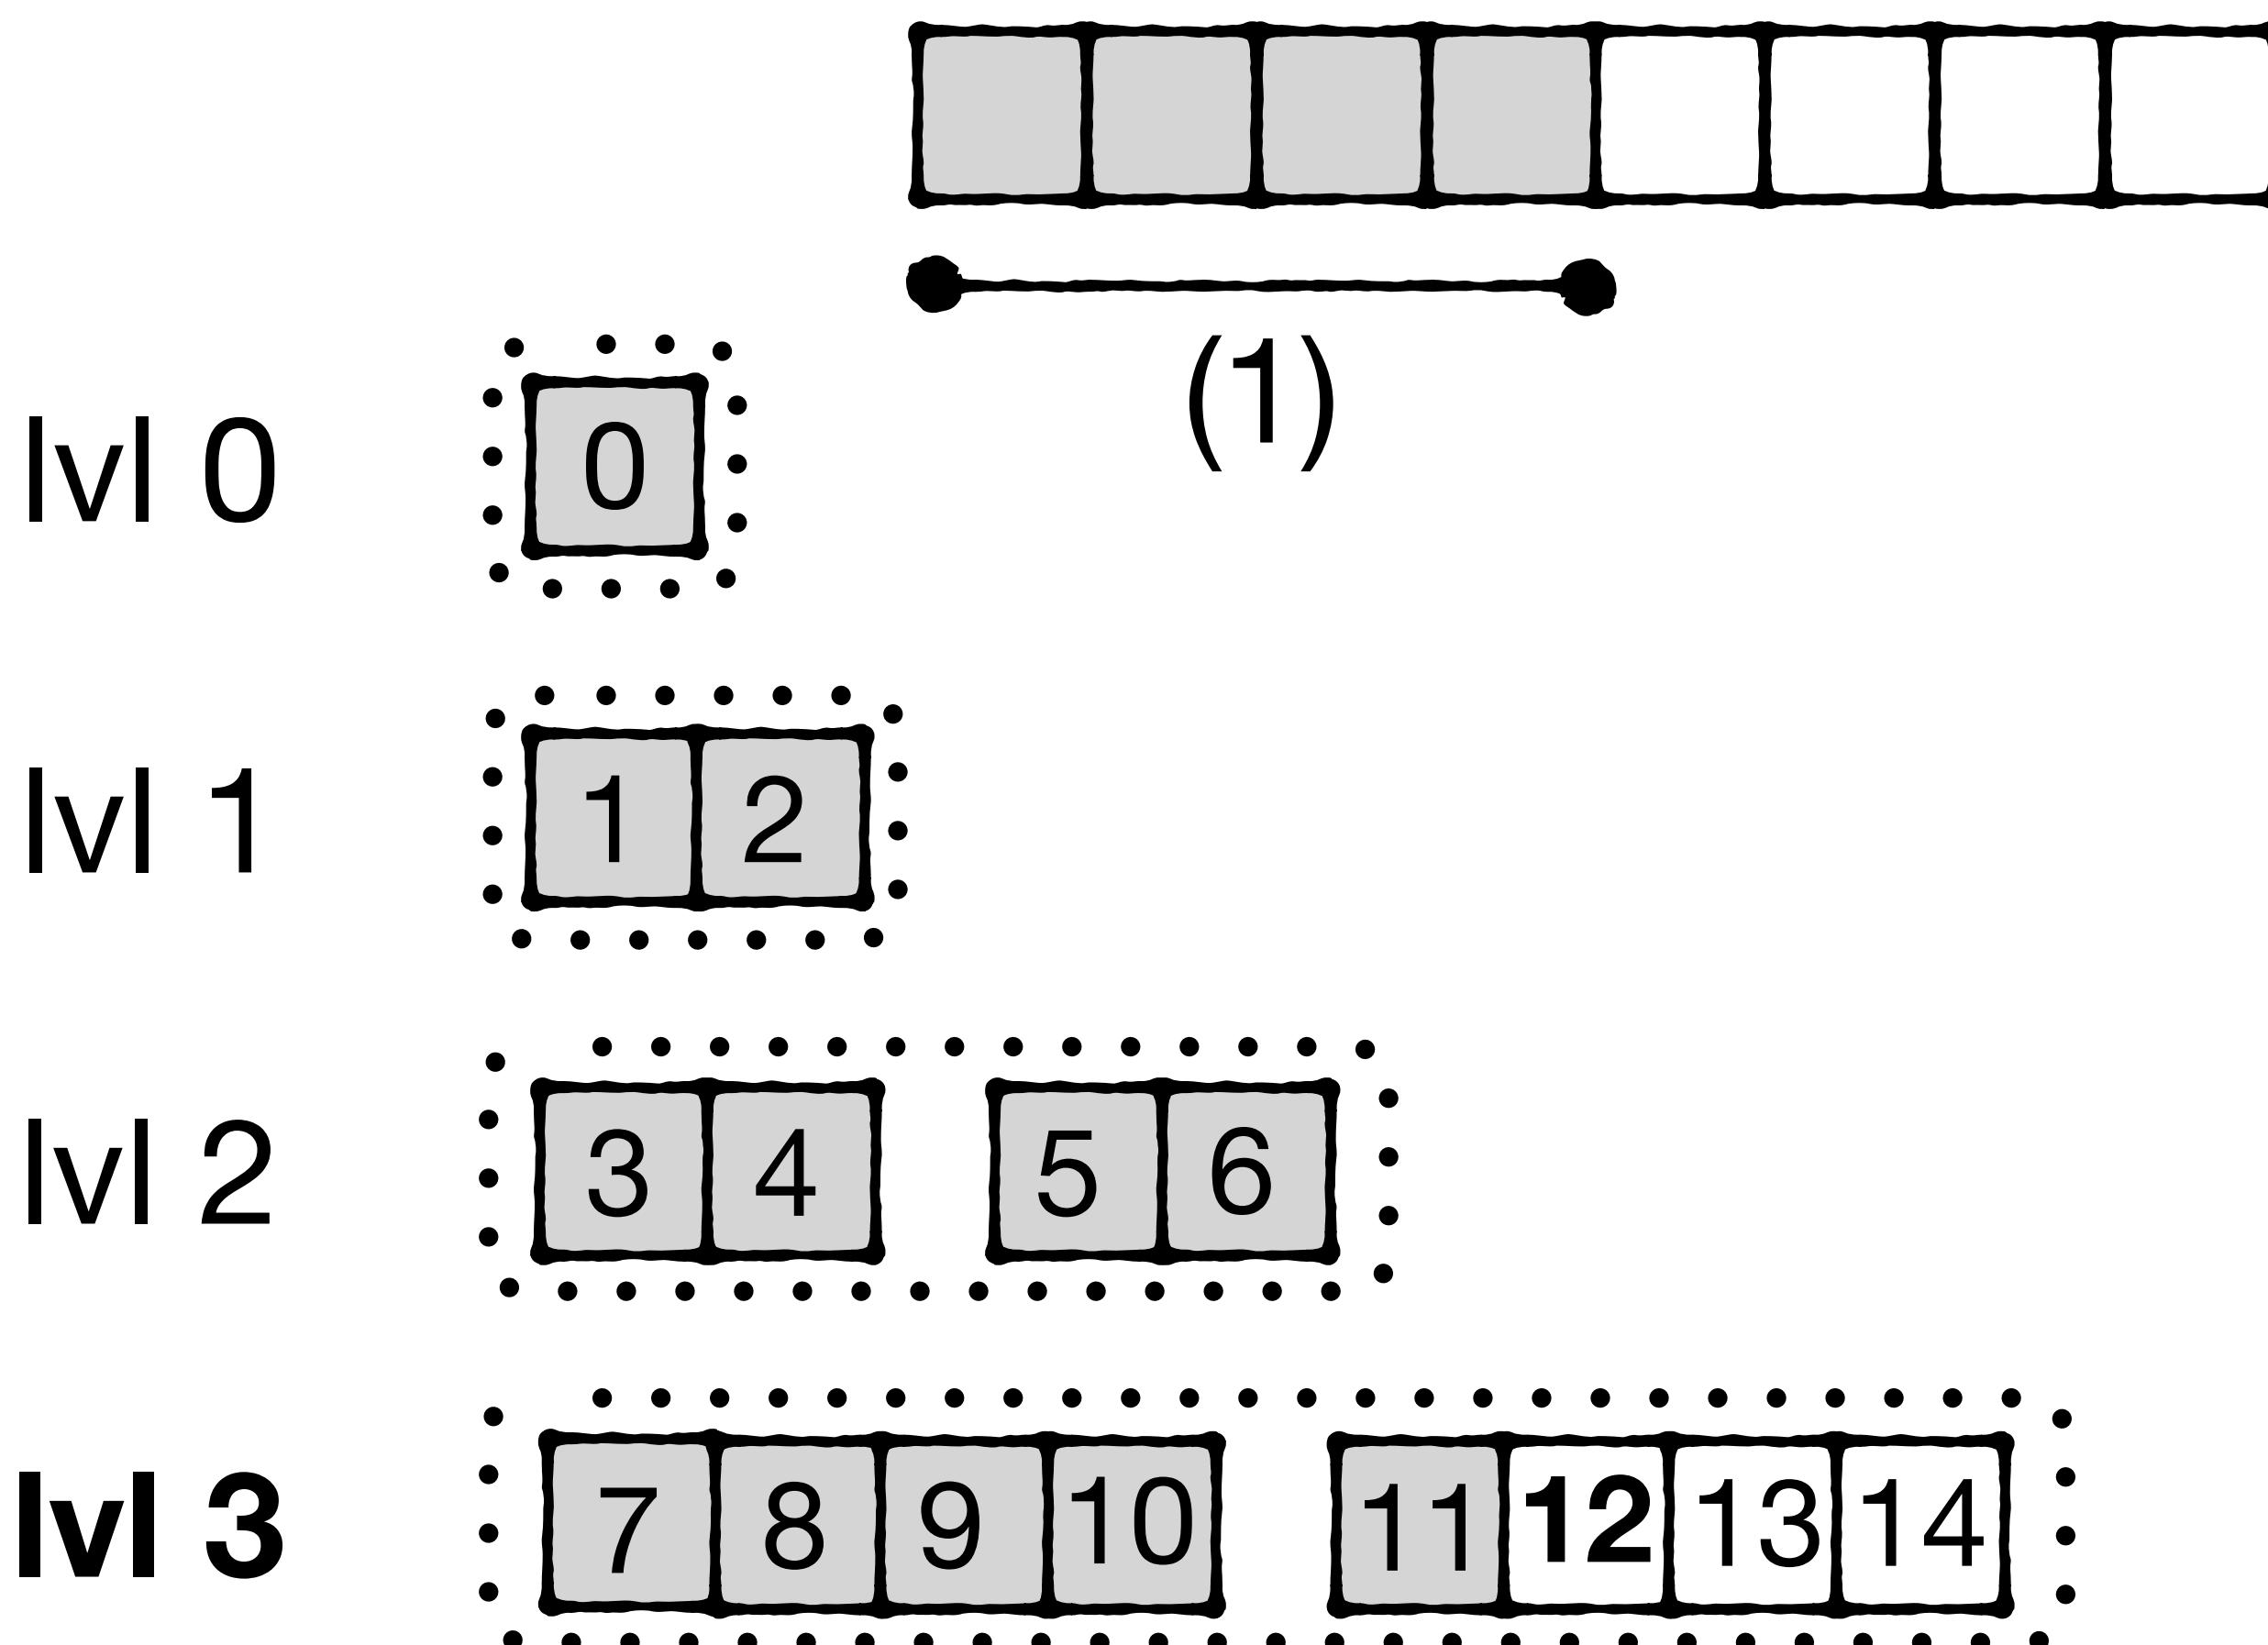


blocks in levels 0 to k-1

$$= 2^{\lfloor k/2 \rfloor} \cdot (2 + (k \bmod 2)) - 2$$

Intuition: number of new data blocks grows every other level

get(i)

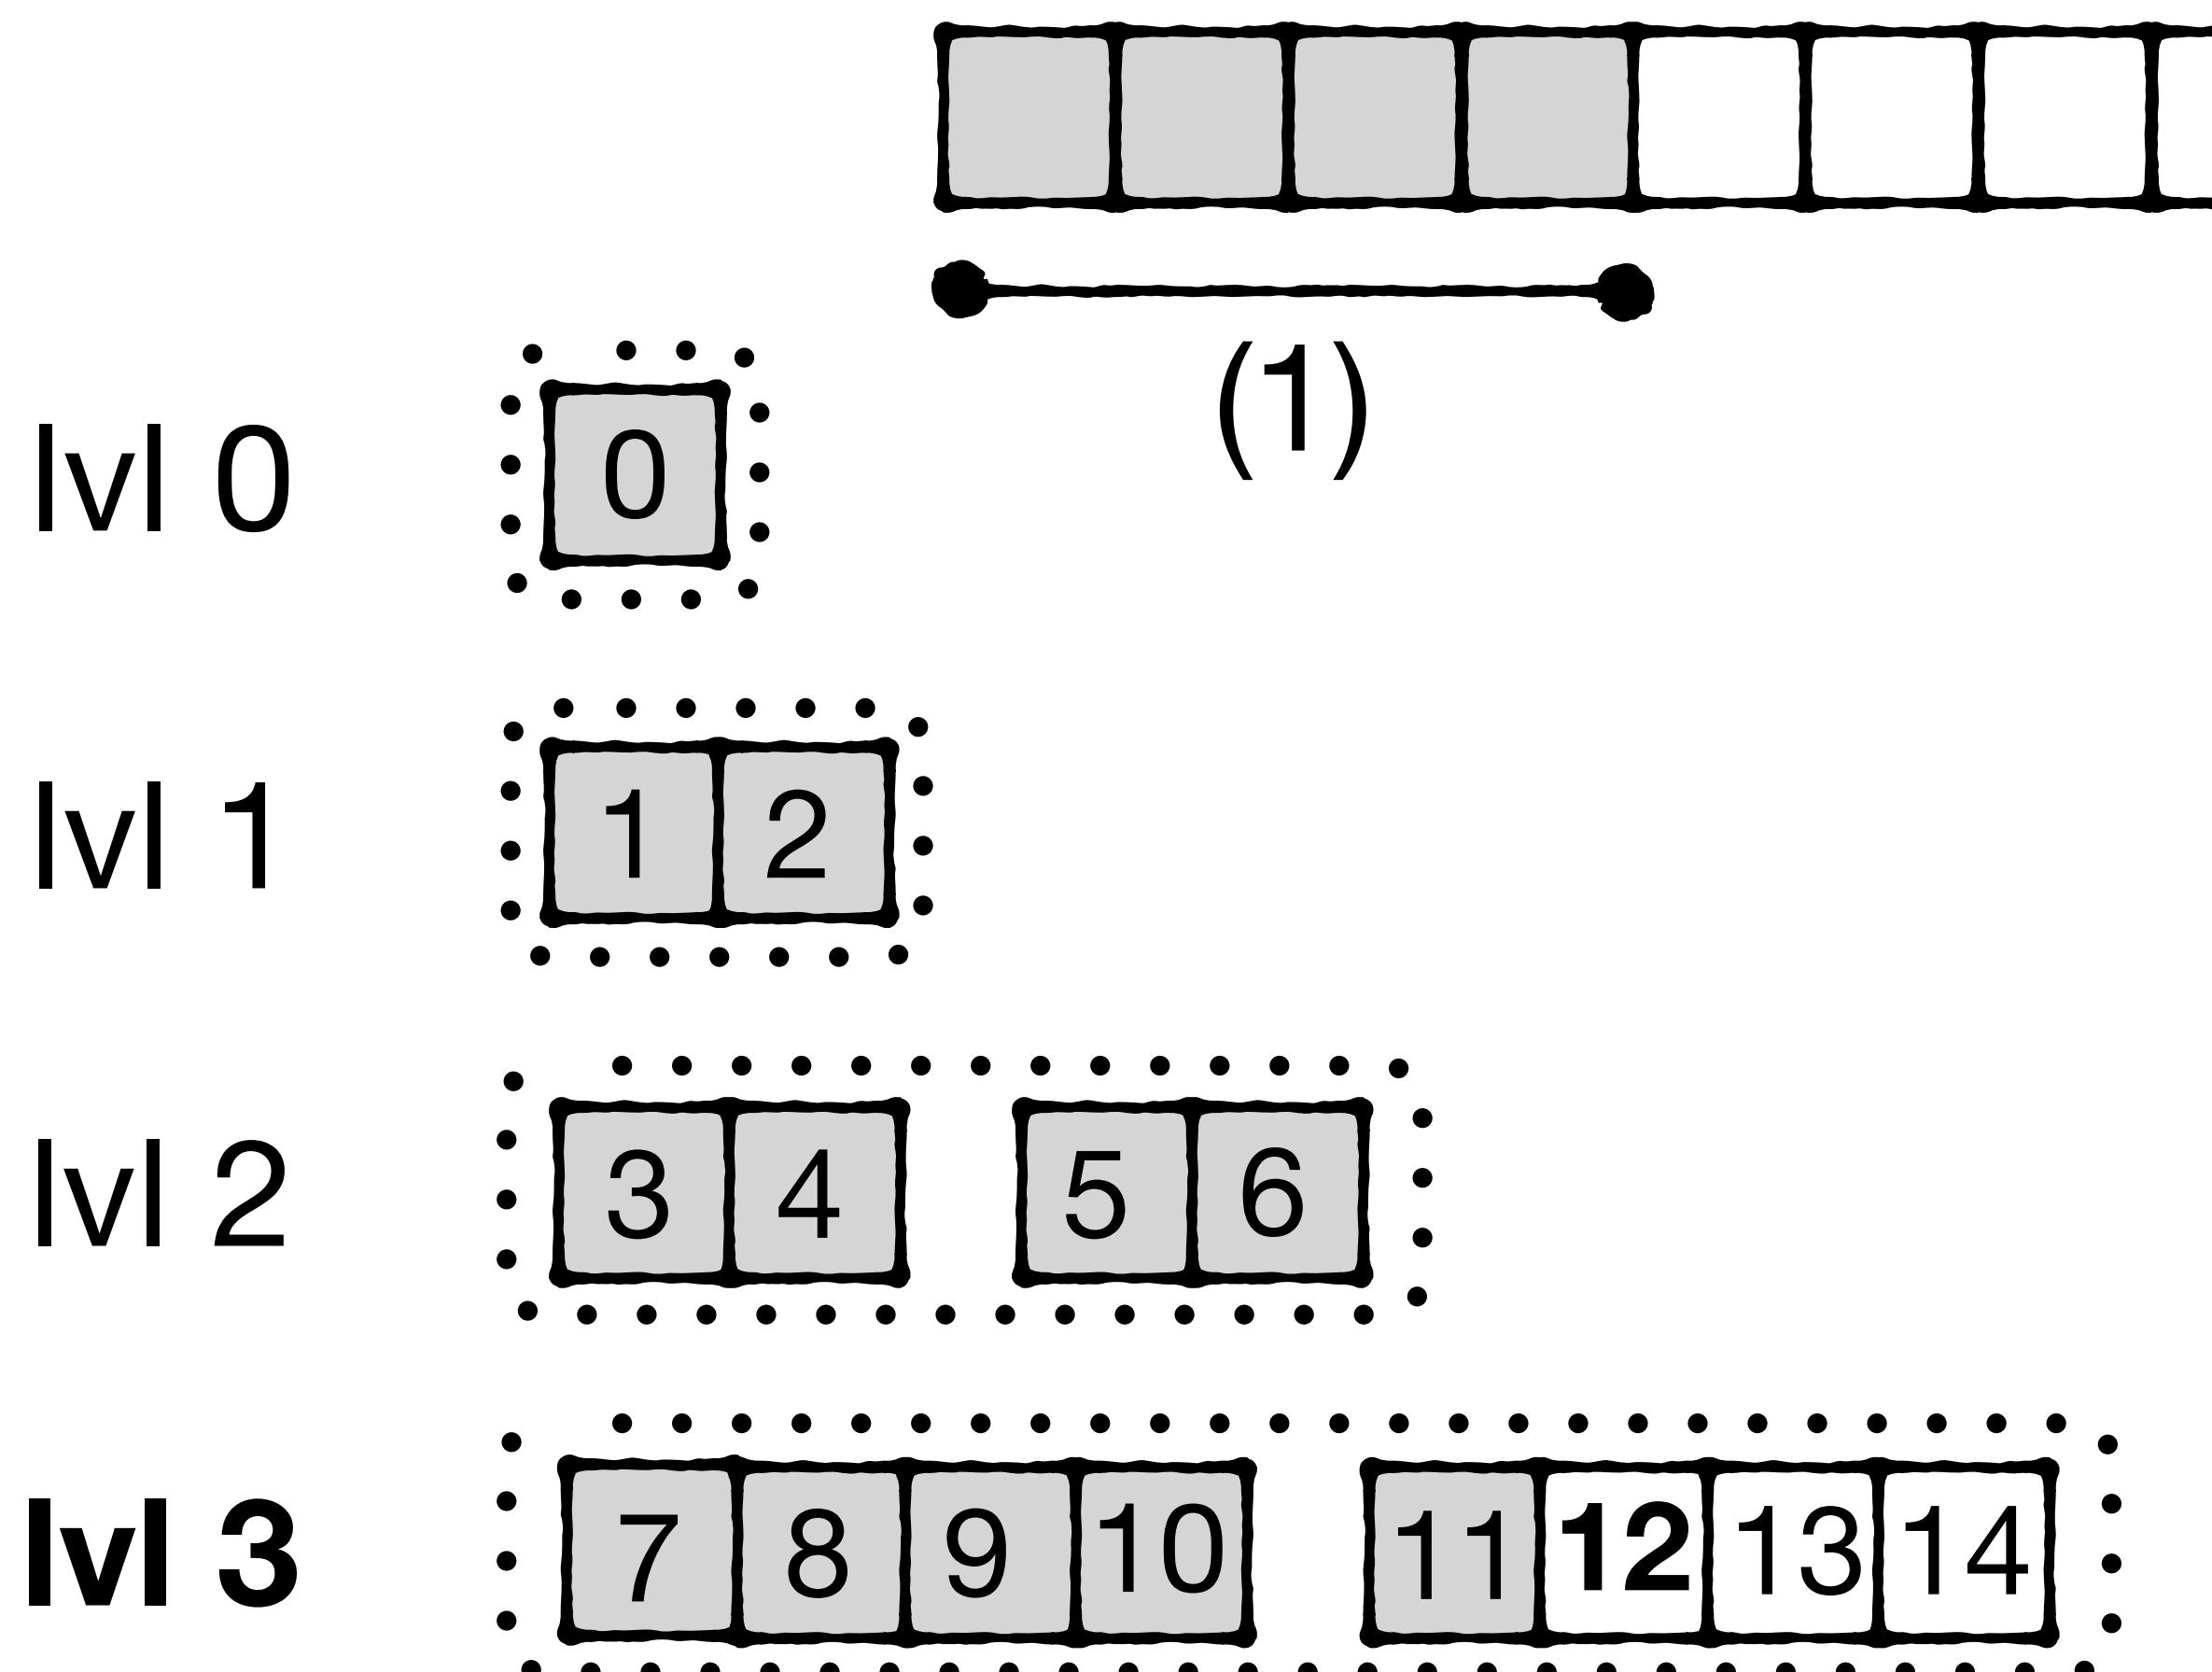


blocks in levels 0 to k-1

$$= 2^{\lfloor k/2 \rfloor} \cdot (2 + (k \bmod 2)) - 2$$

Level k	# Blocks
0	0
1	1
2	2
3	4
4	6
5	10
6	14
7	22
...	...

get(i)

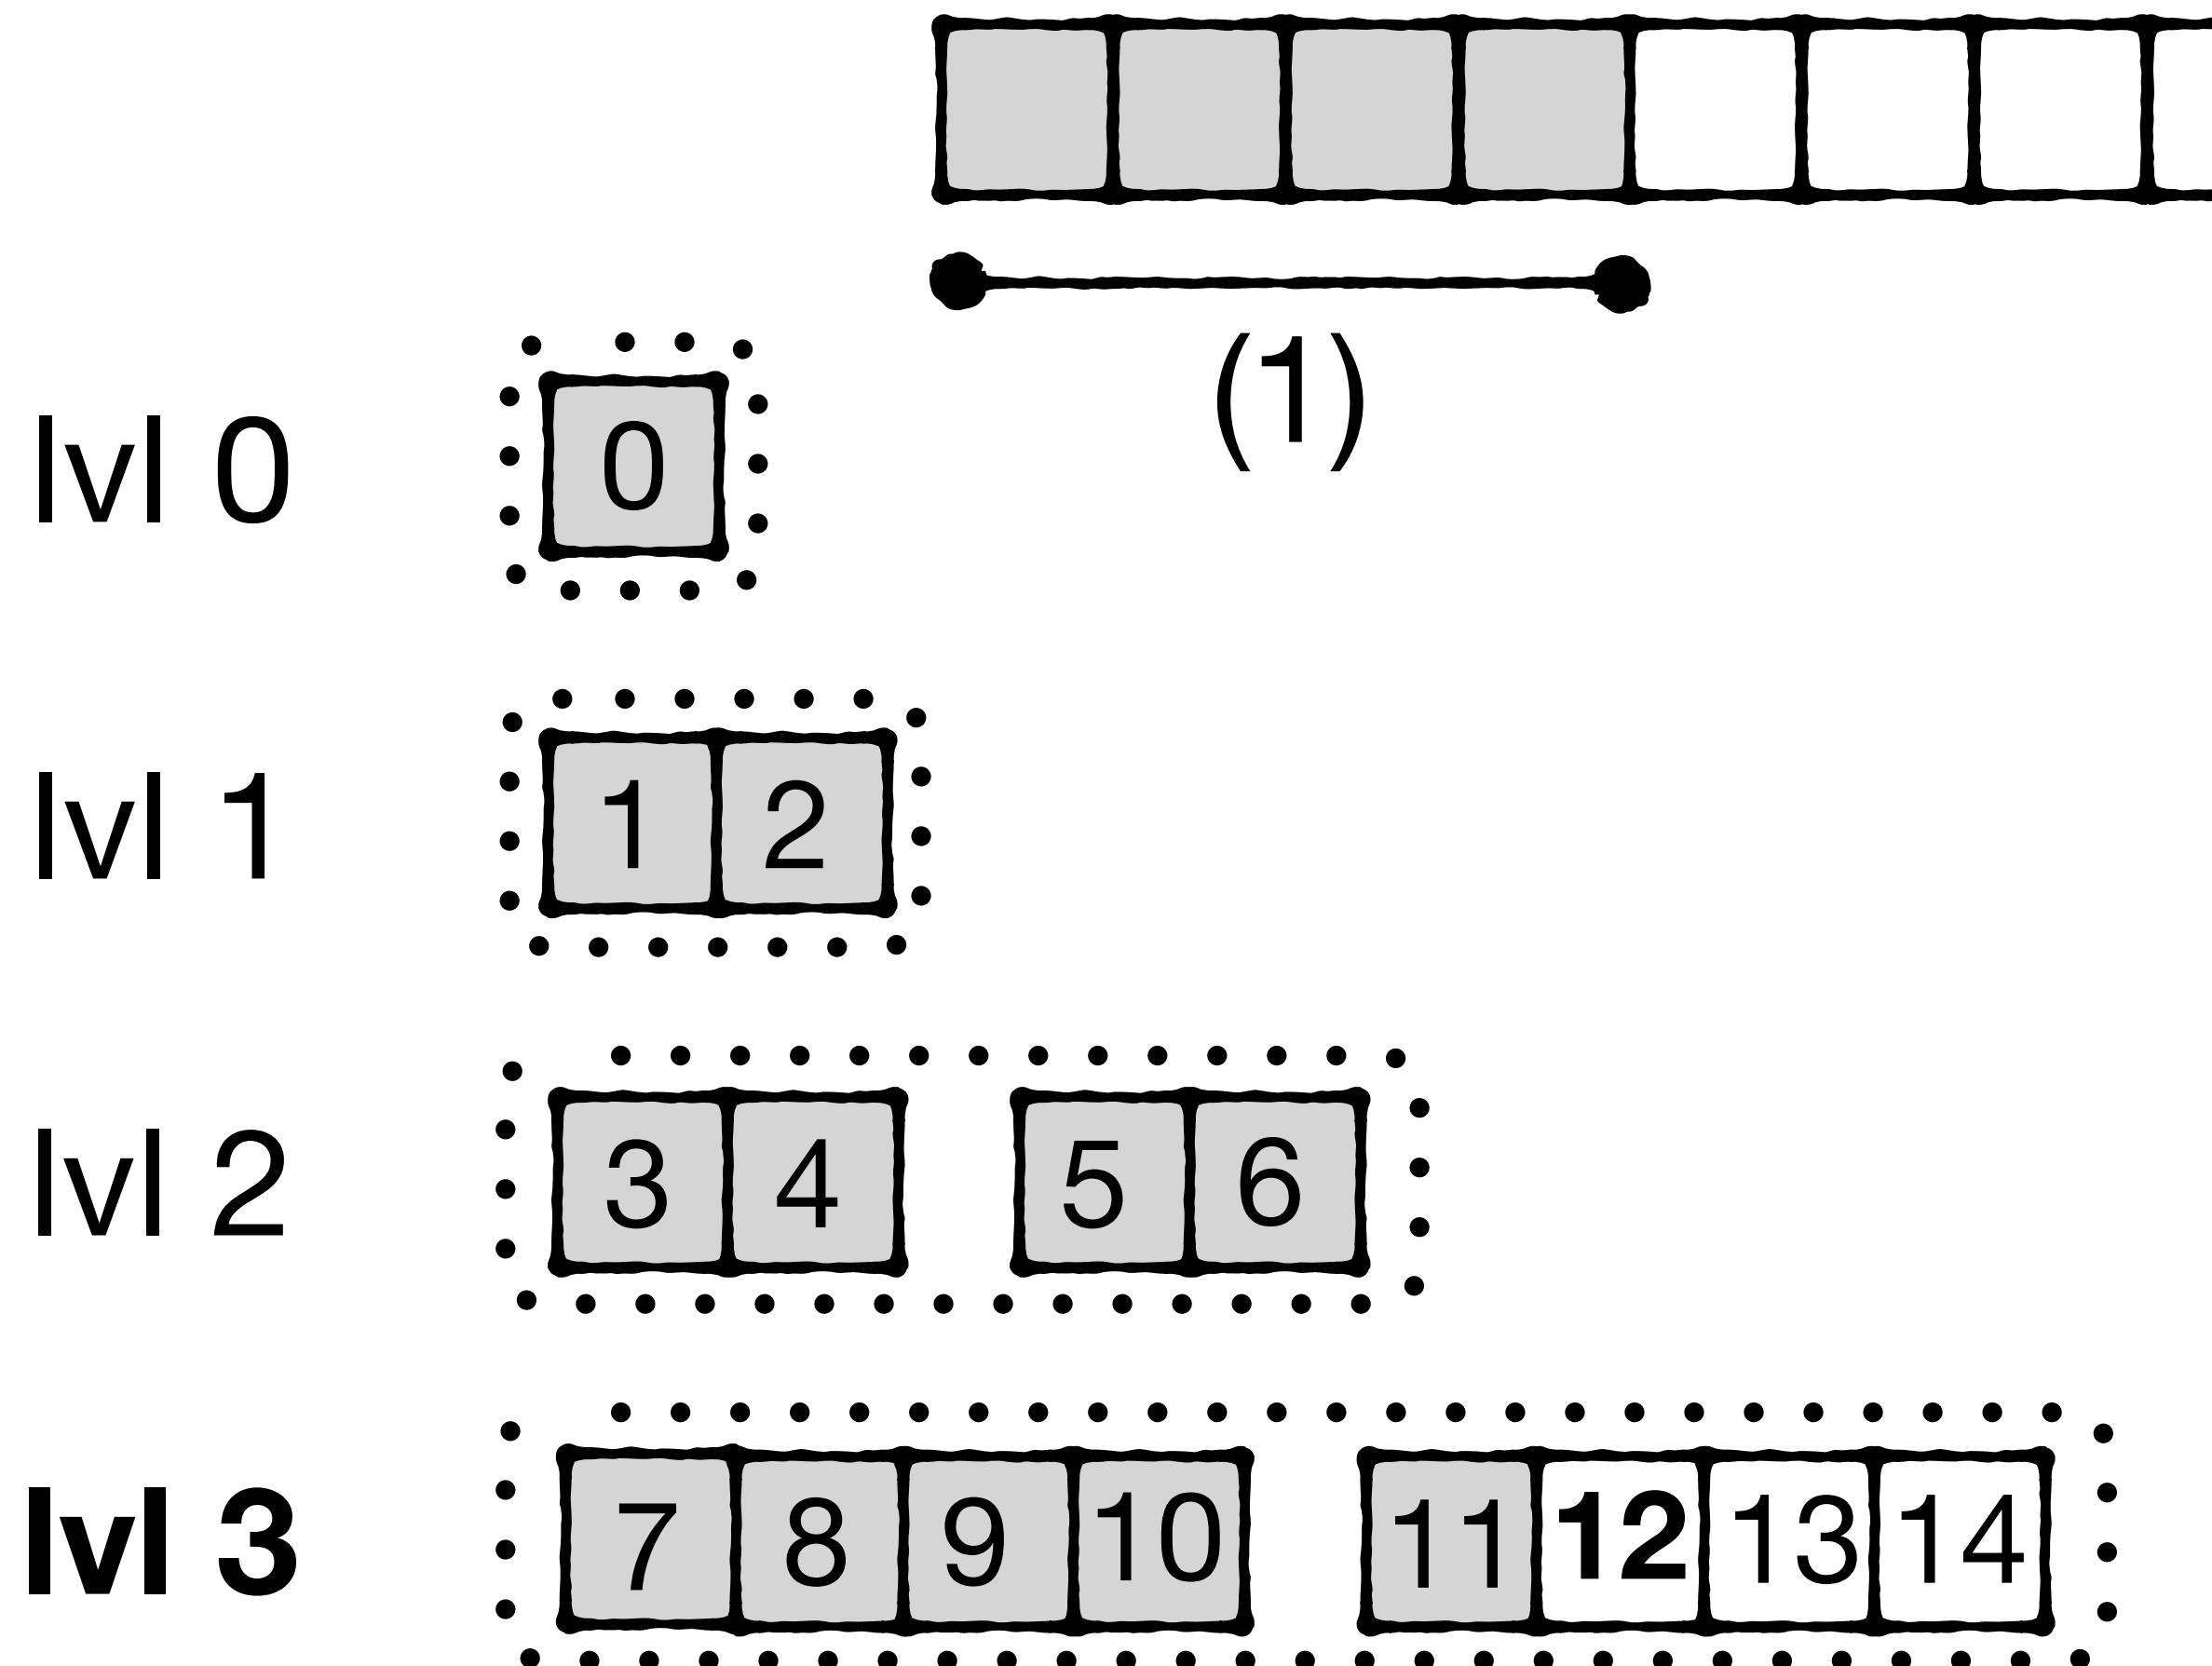


blocks in levels 0 to k-1

$$= 2^{\lfloor k/2 \rfloor} \cdot (2 + (k \bmod 2)) - 2$$

Integer division is slow

get(i)



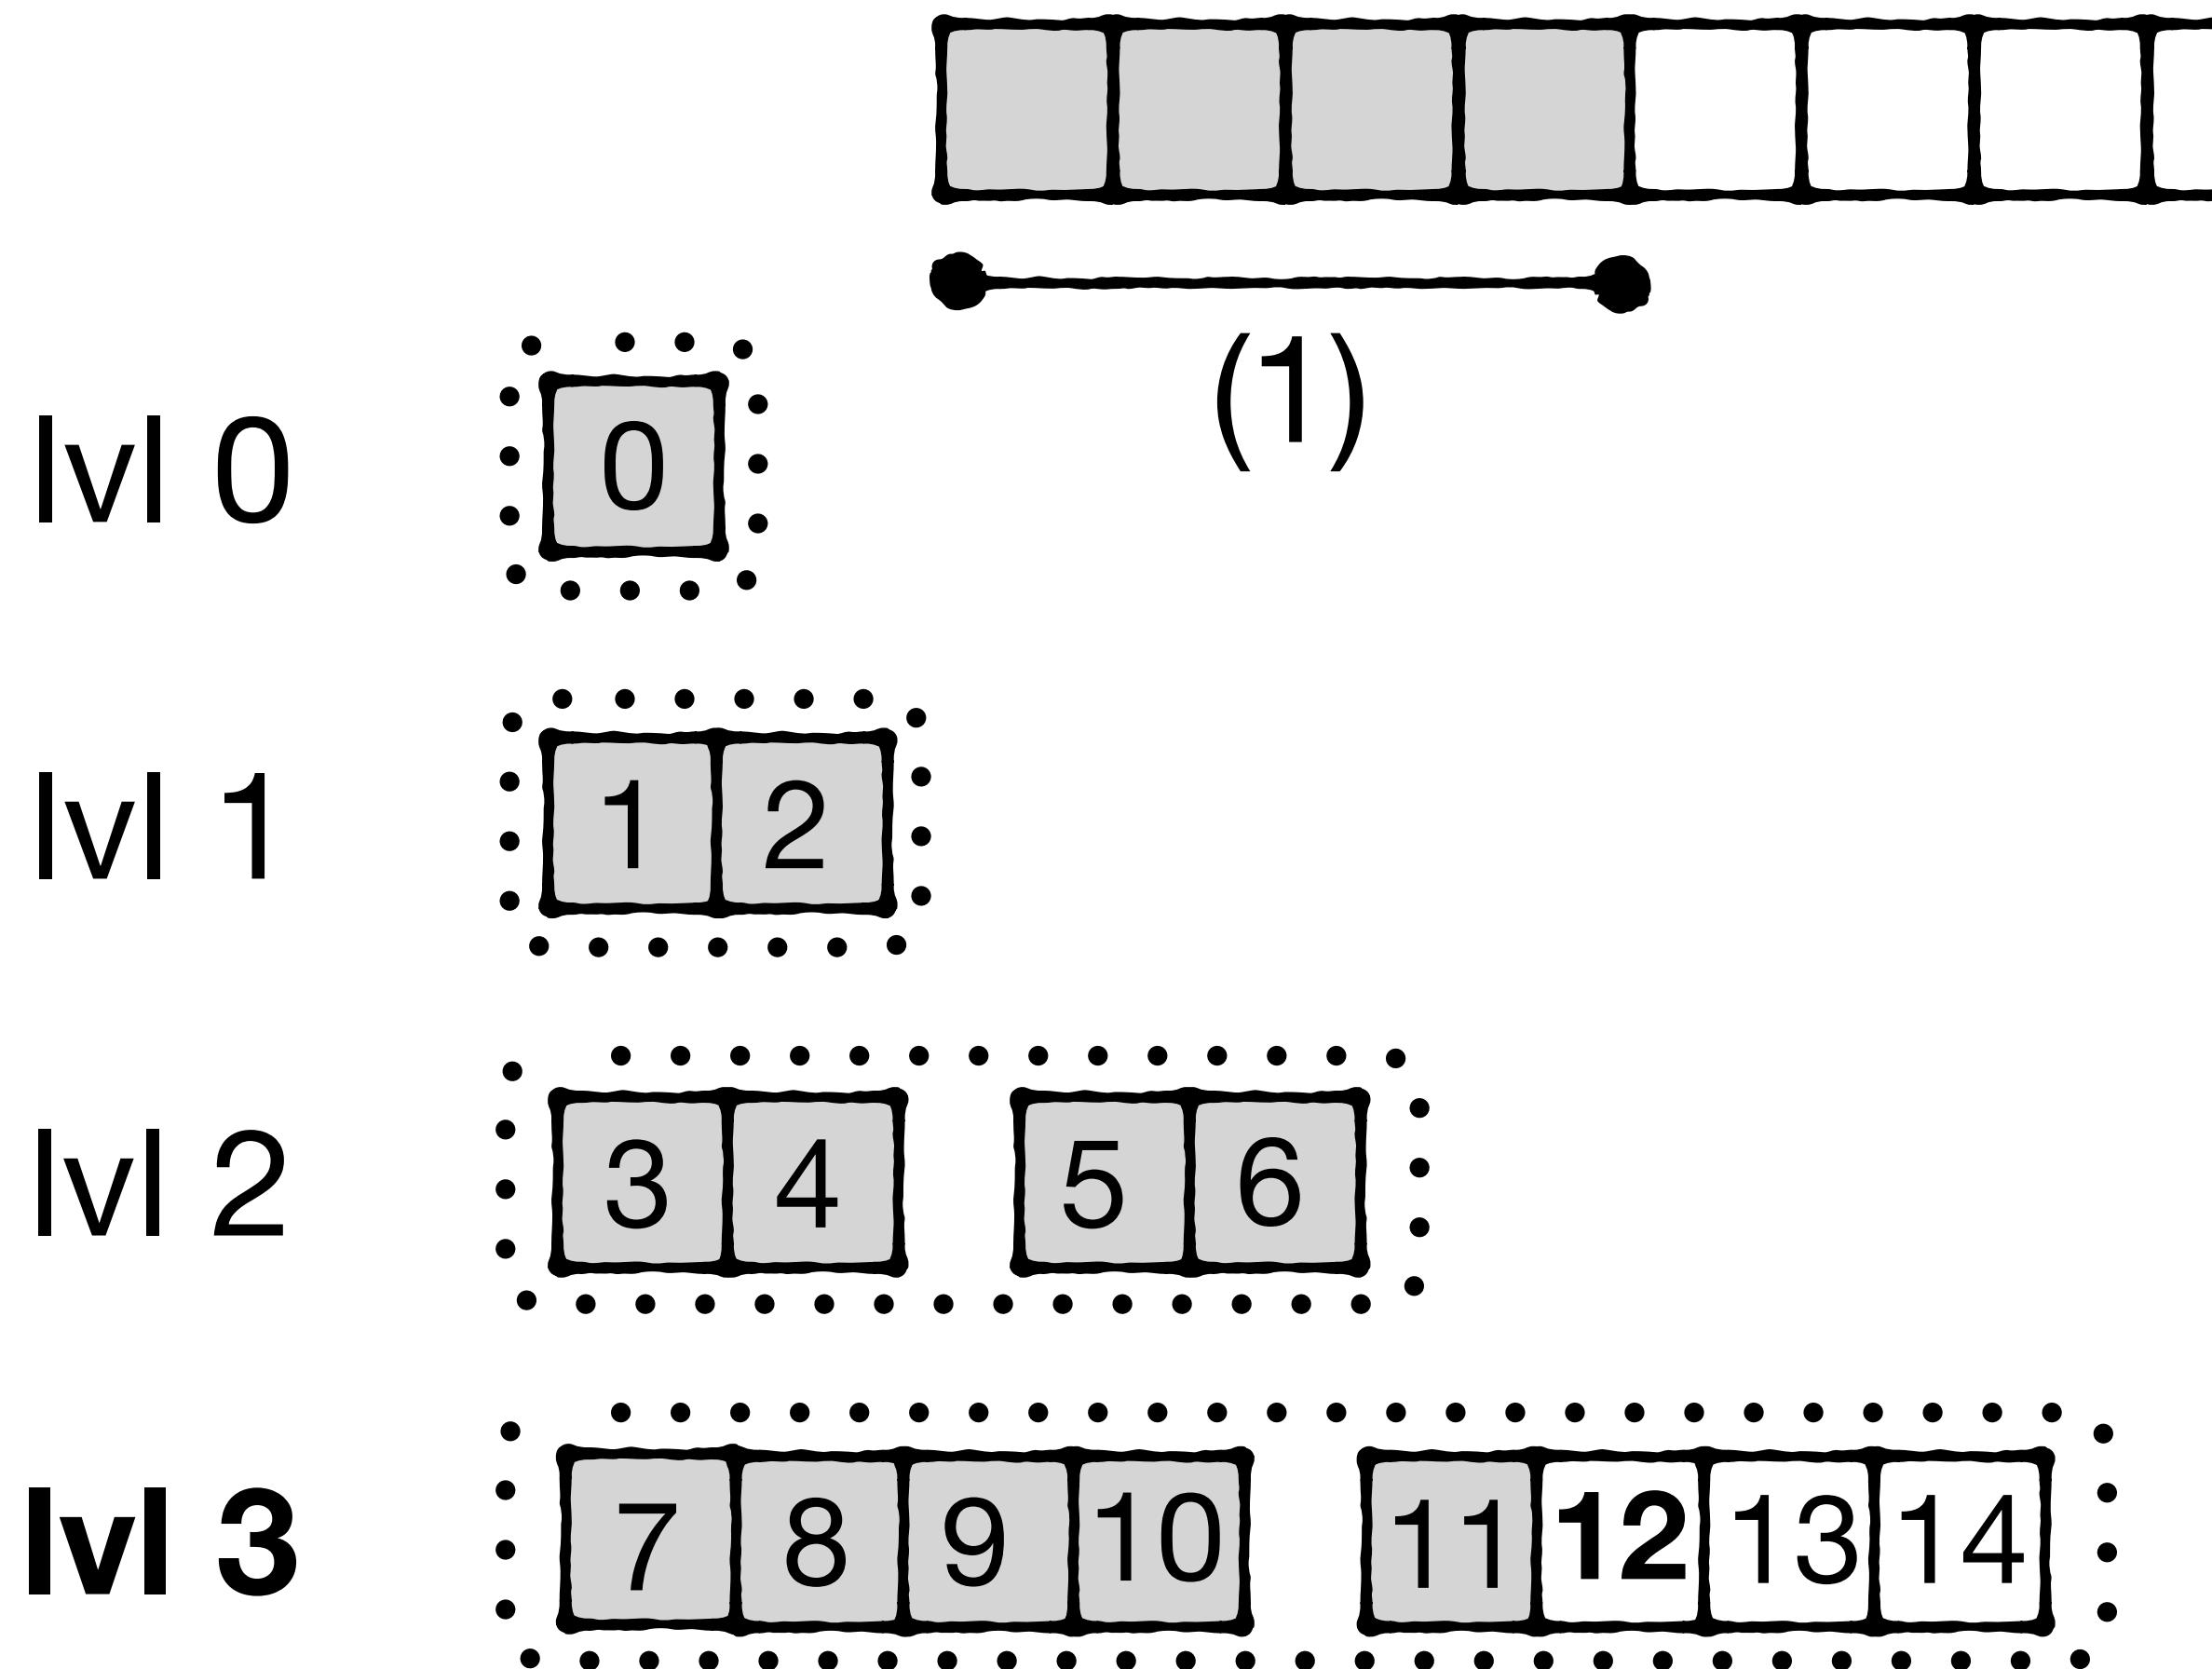
blocks in levels 0 to k-1

$$= 2^{\lfloor k/2 \rfloor} \cdot (2 + (k \bmod 2)) - 2$$



Power is slow

get(i)

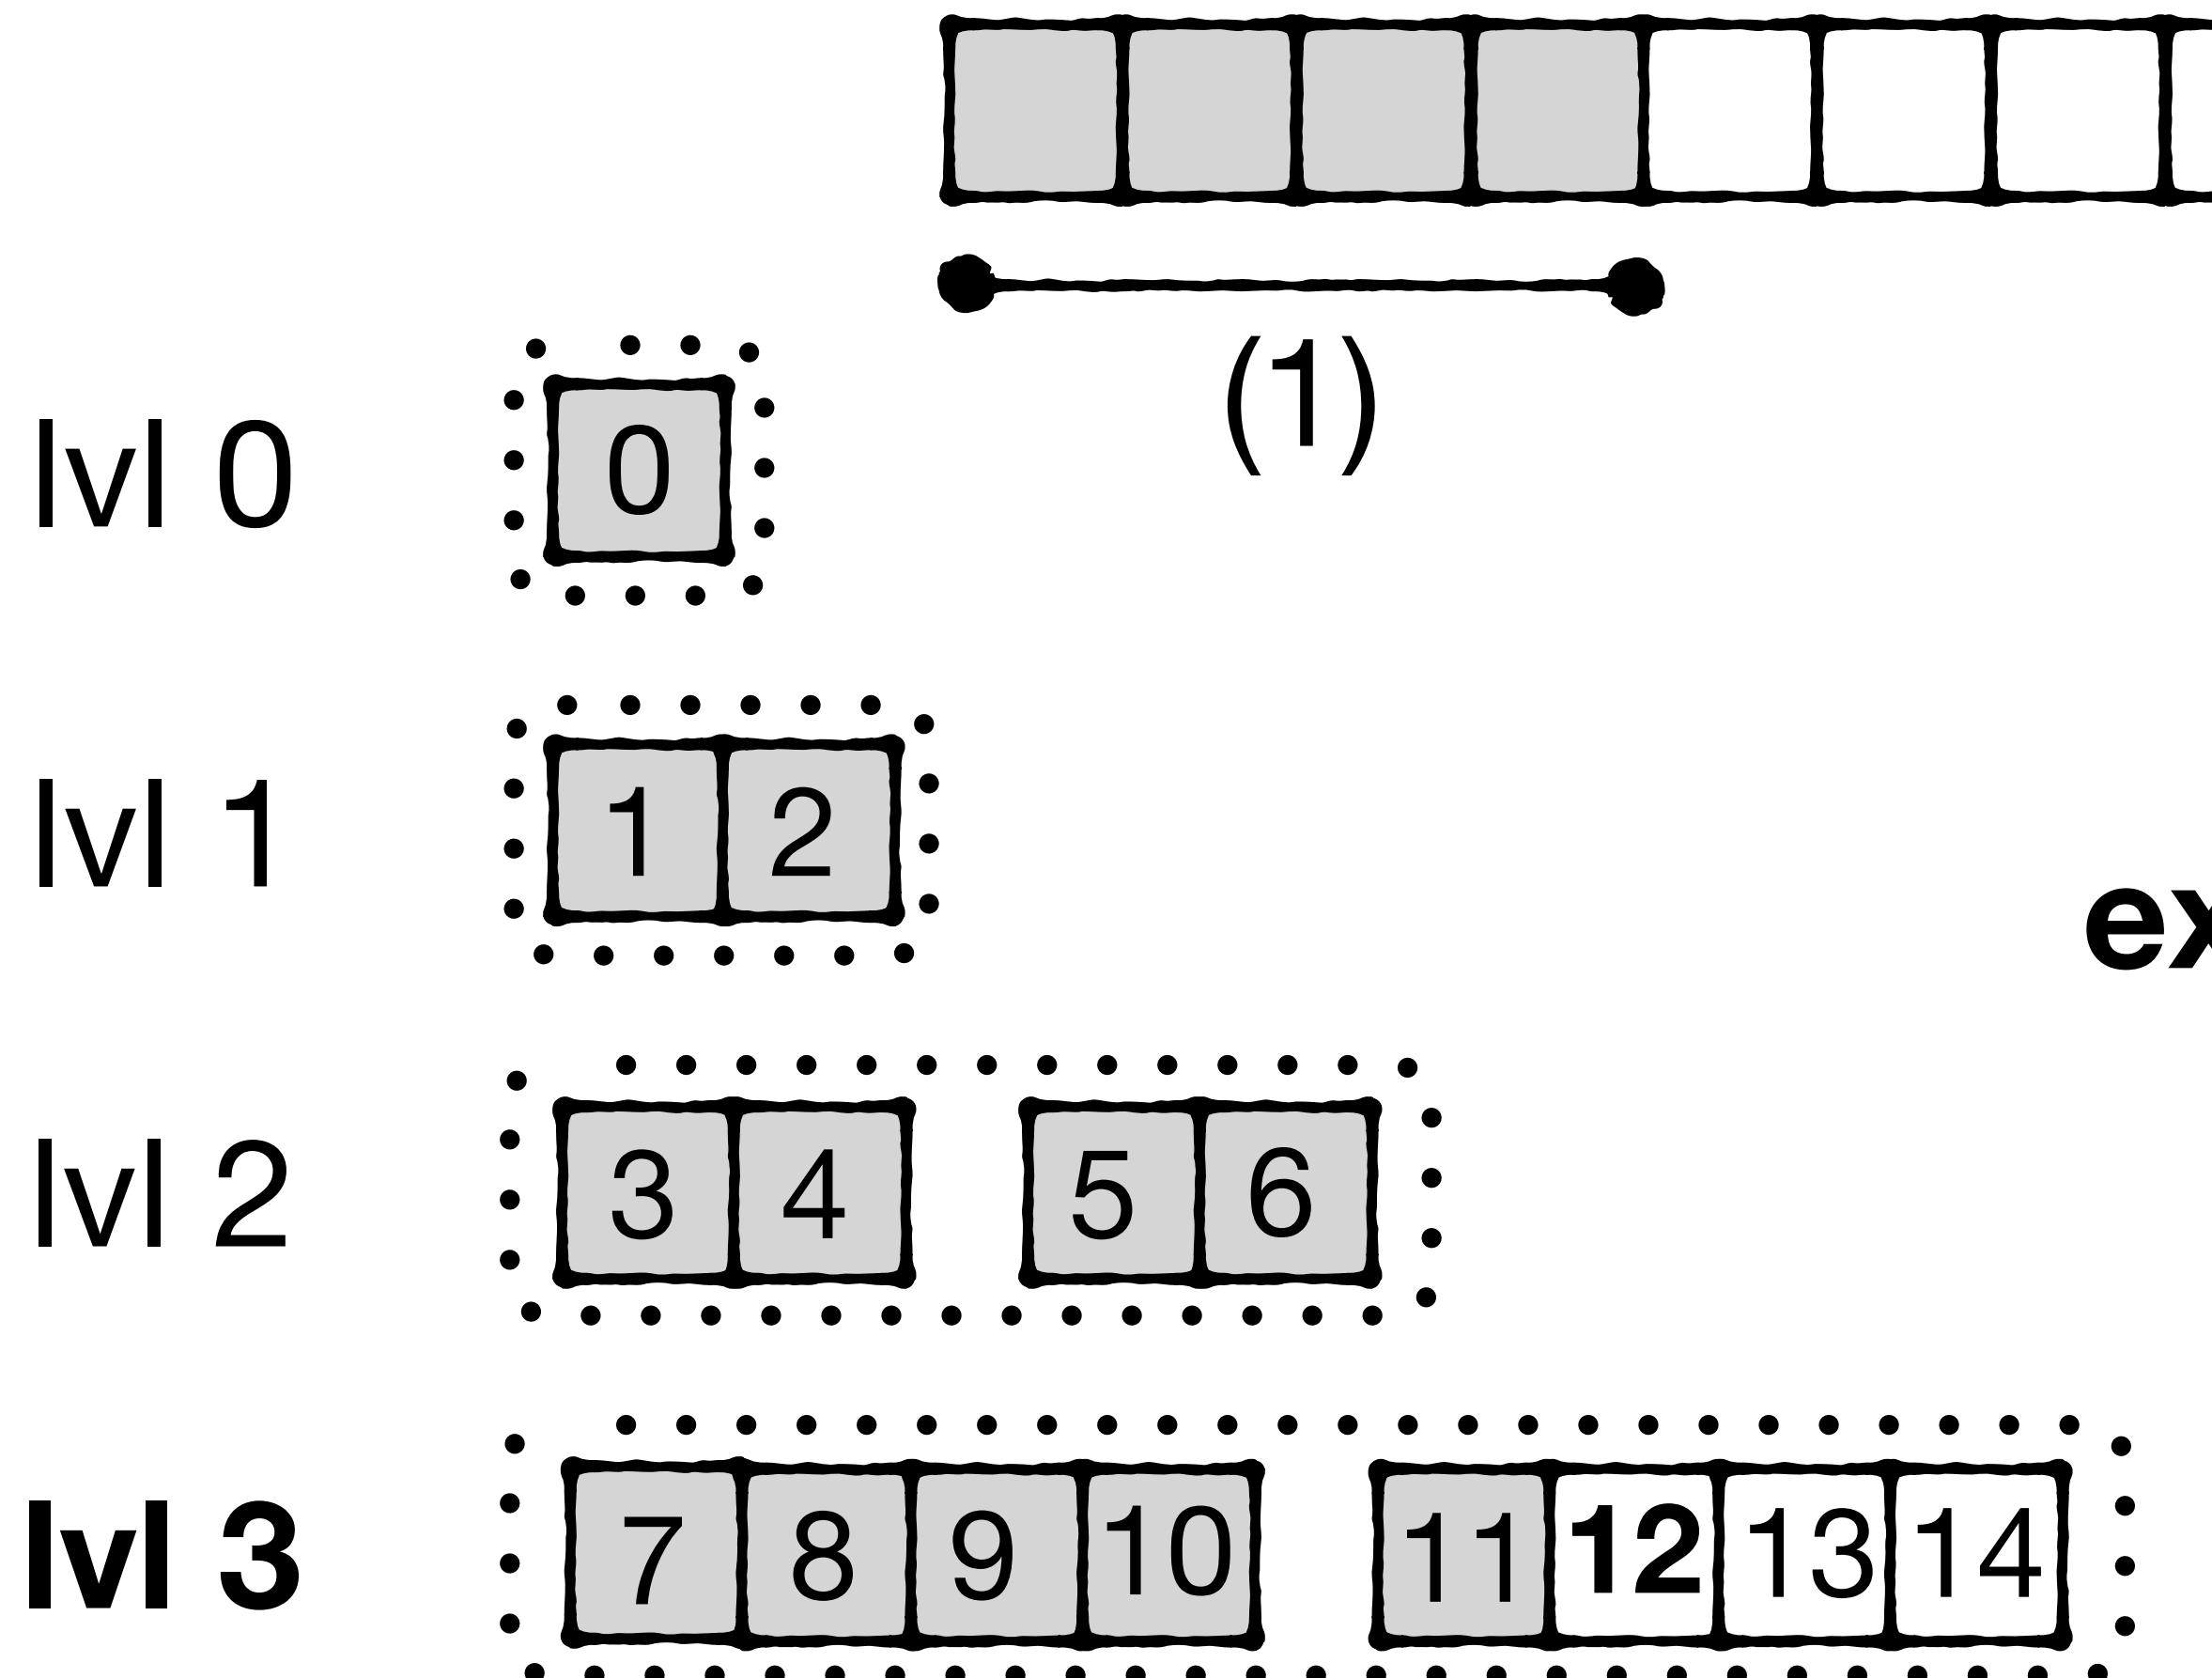


blocks in levels 0 to k-1

$$= 2^{\lfloor k/2 \rfloor} \cdot (2 + (k \bmod 2)) - 2$$

How to speed up?

get(i)

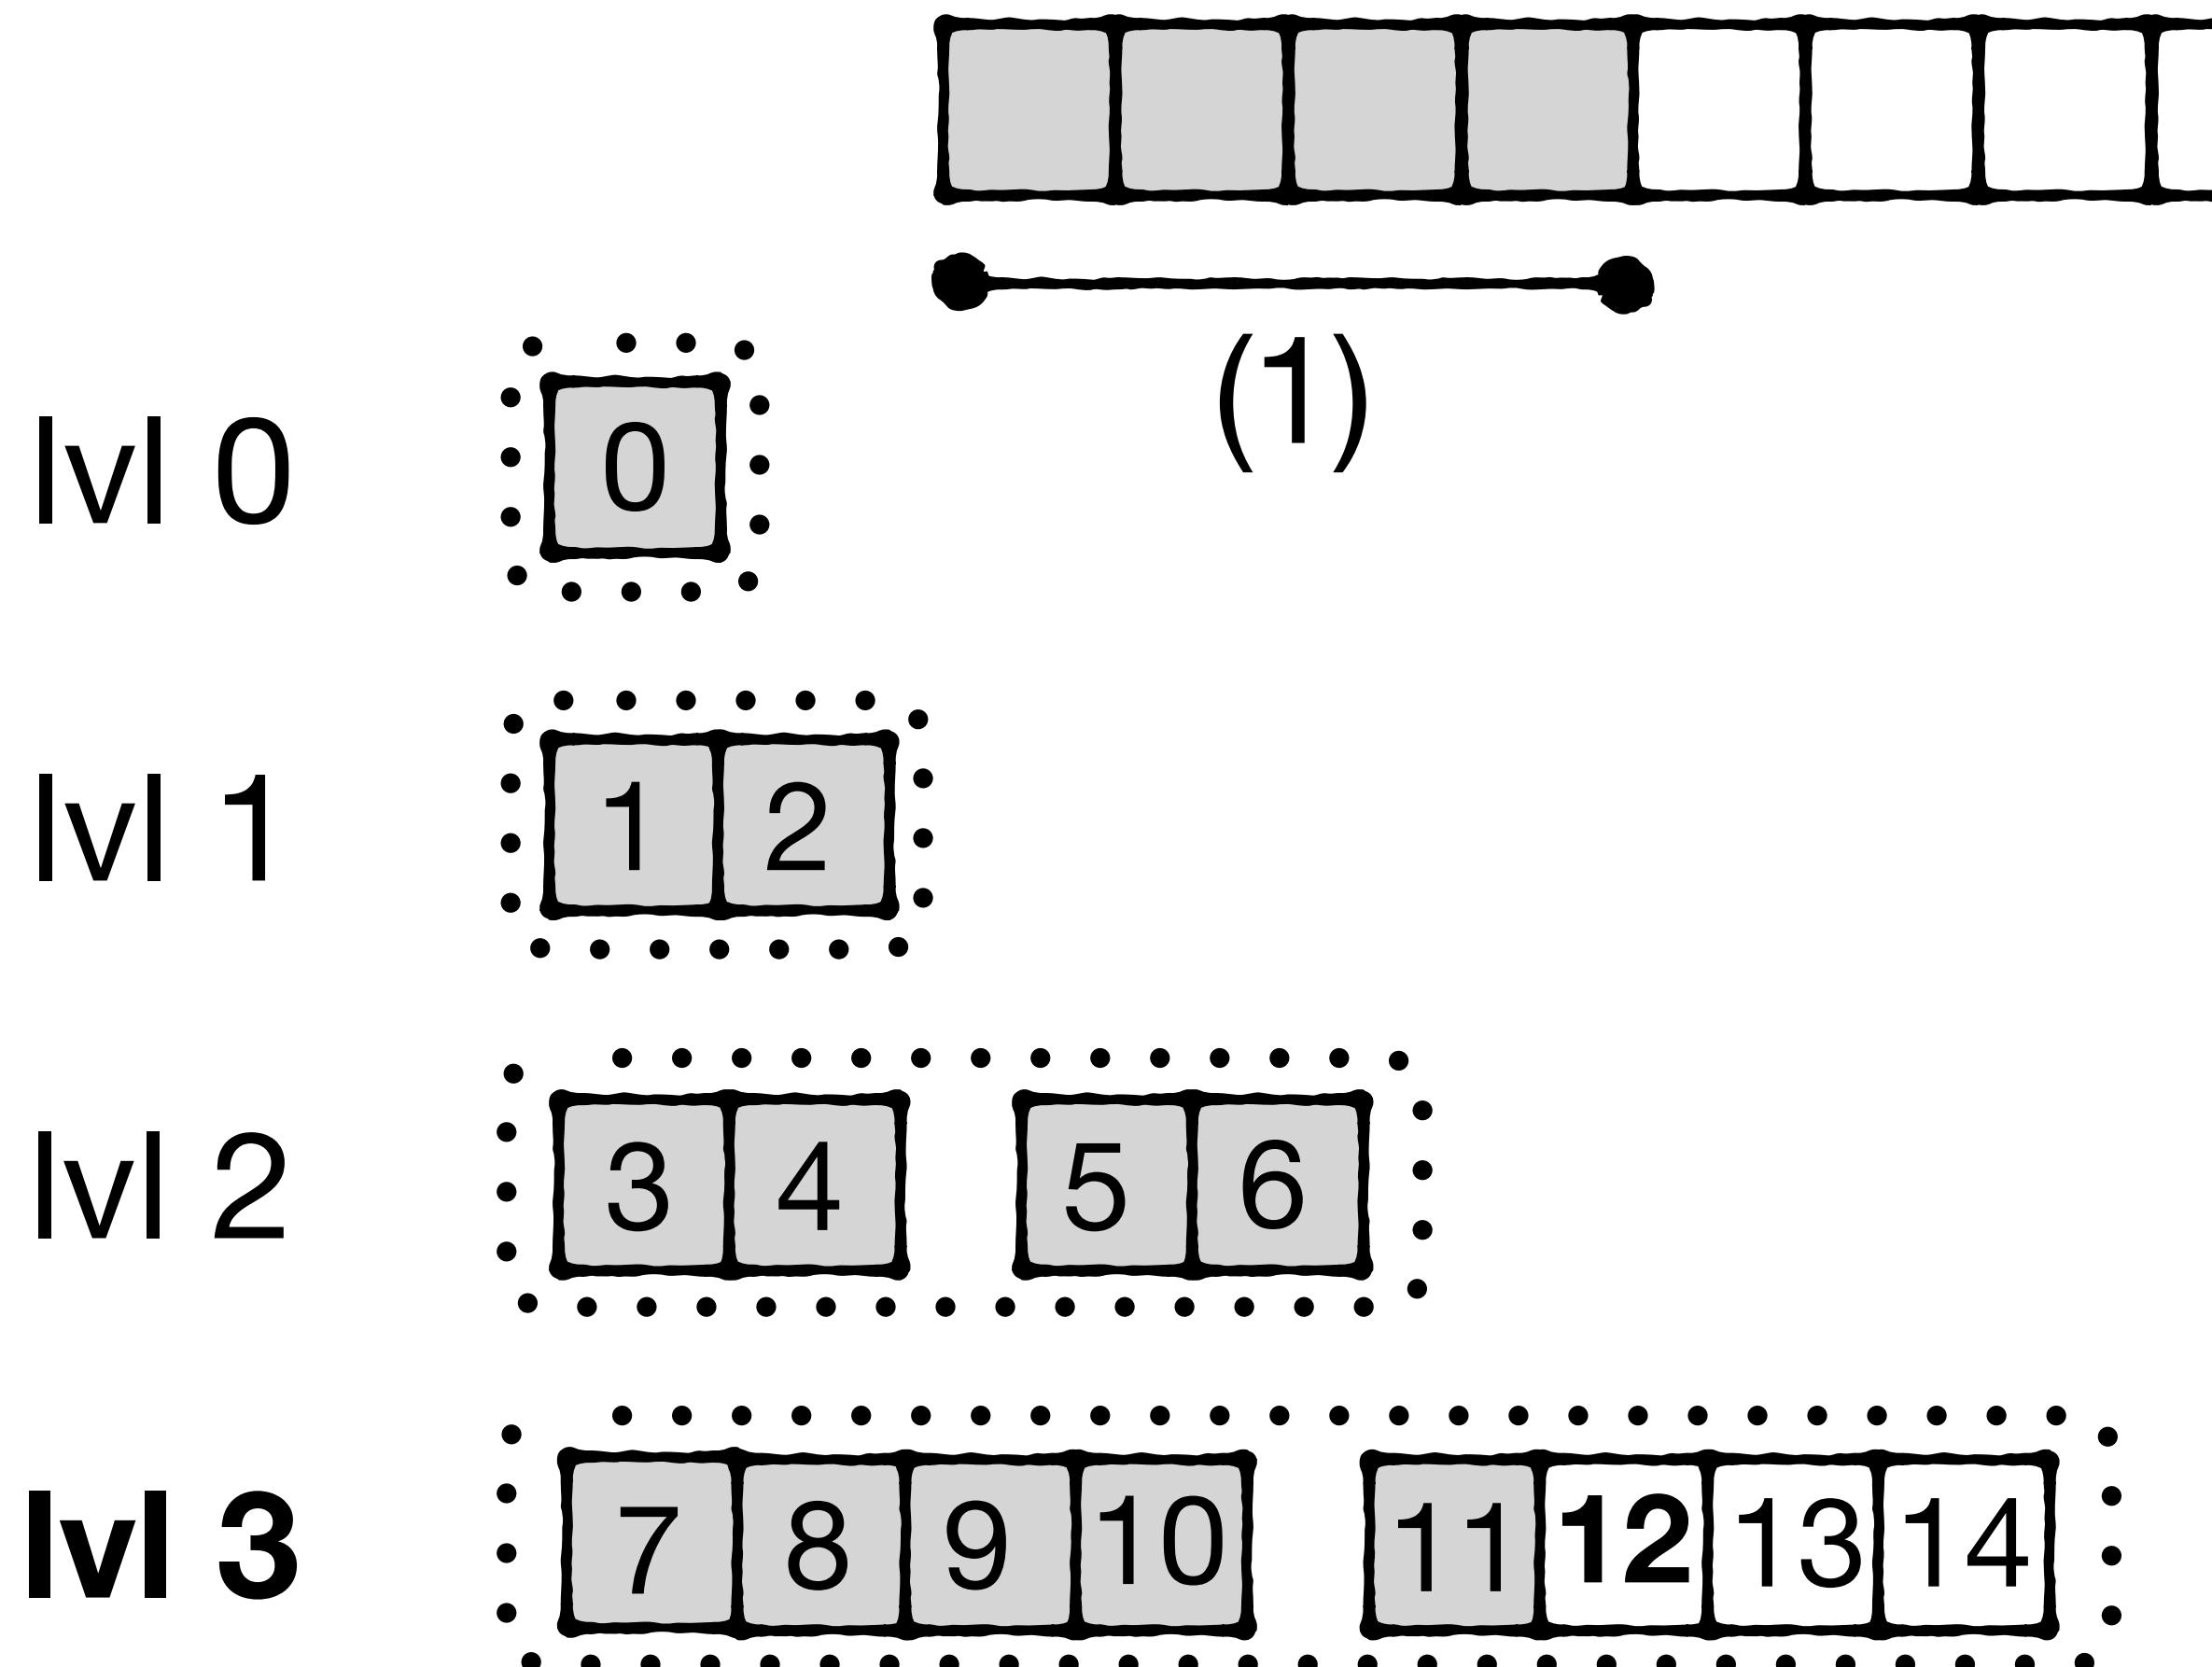


blocks in levels 0 to k-1

$$= 2^{\lfloor k/2 \rfloor} \cdot (2 + (k \bmod 2)) - 2$$

Insight: division & exponentiation by 2 can be done with bitwise operators

get(i)



blocks in levels 0 to k-1

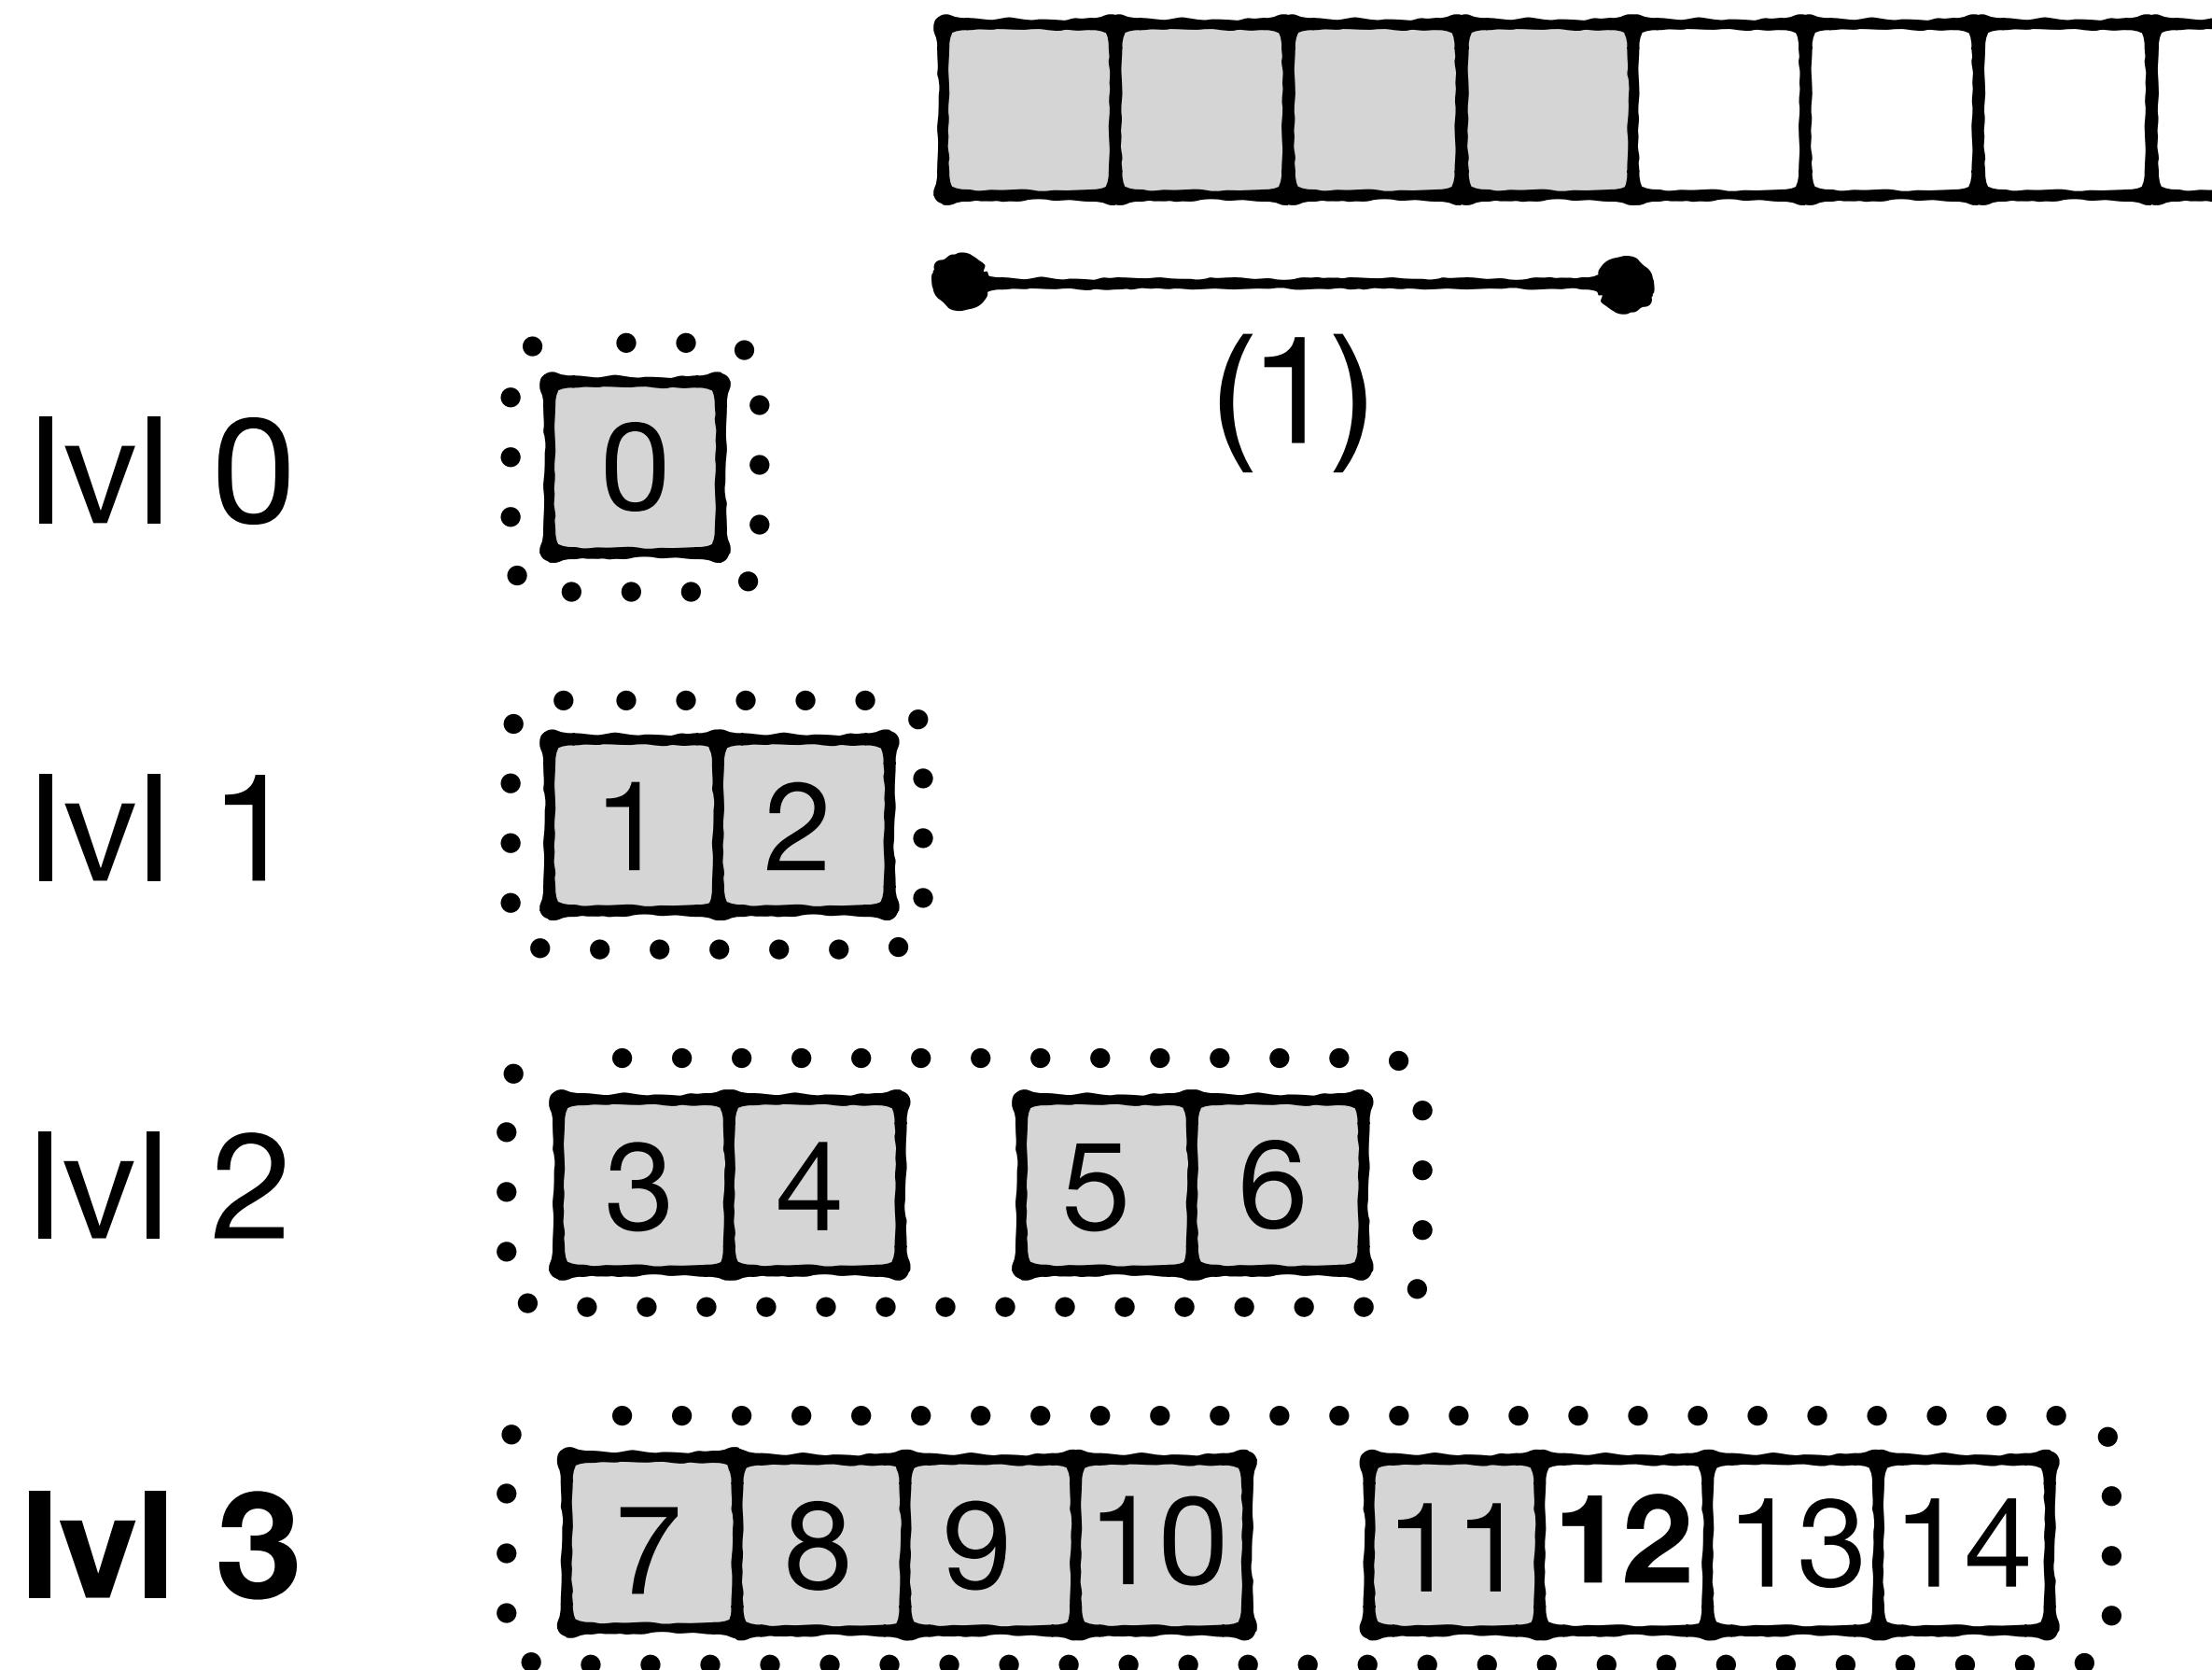
$$= 2^{\lfloor k/2 \rfloor} \cdot (2 + (k \bmod 2)) - 2$$

$$= 2^{\lfloor k/2 \rfloor} \cdot (2 + (k \& 1)) - 2$$



“and” with 1

get(i)



blocks in levels 0 to k-1

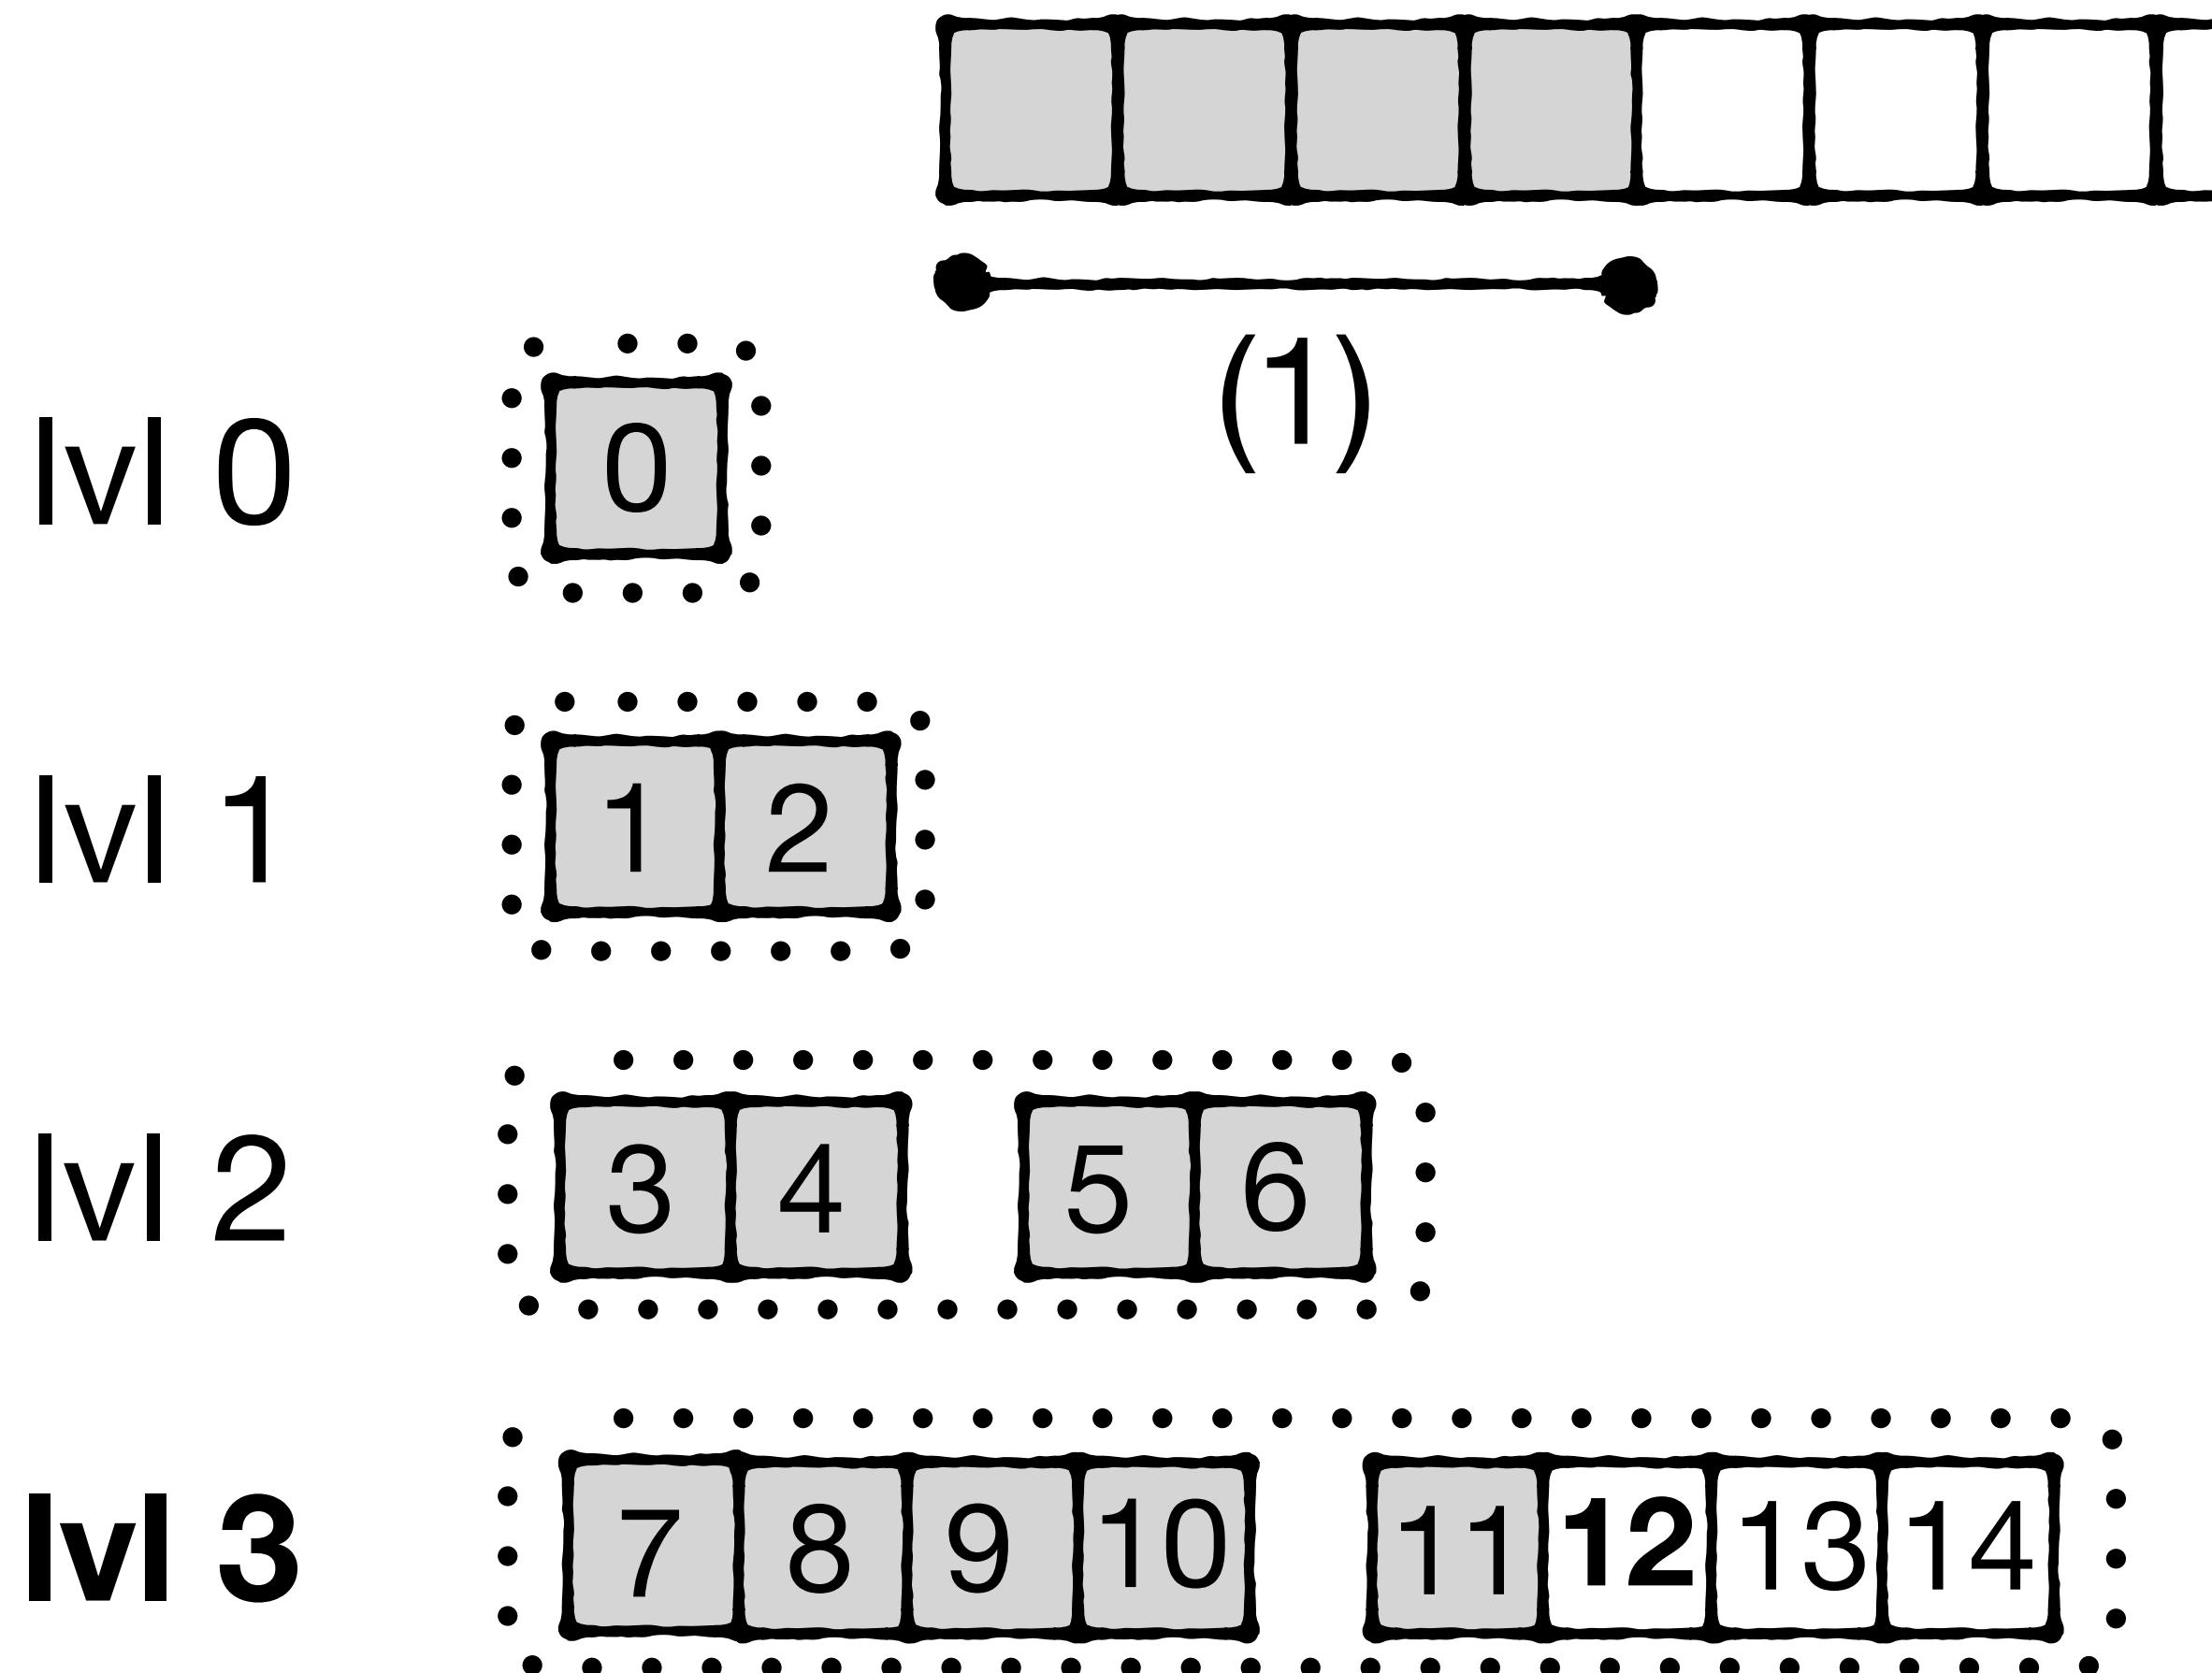
$$= 2^{\lfloor k/2 \rfloor} \cdot (2 + (k \bmod 2)) - 2$$

$$= 2^{(k \gg 1)} \cdot (2 + (k \& 1)) - 2$$



Shift by 1 bit to right

get(i)



blocks in levels 0 to k-1

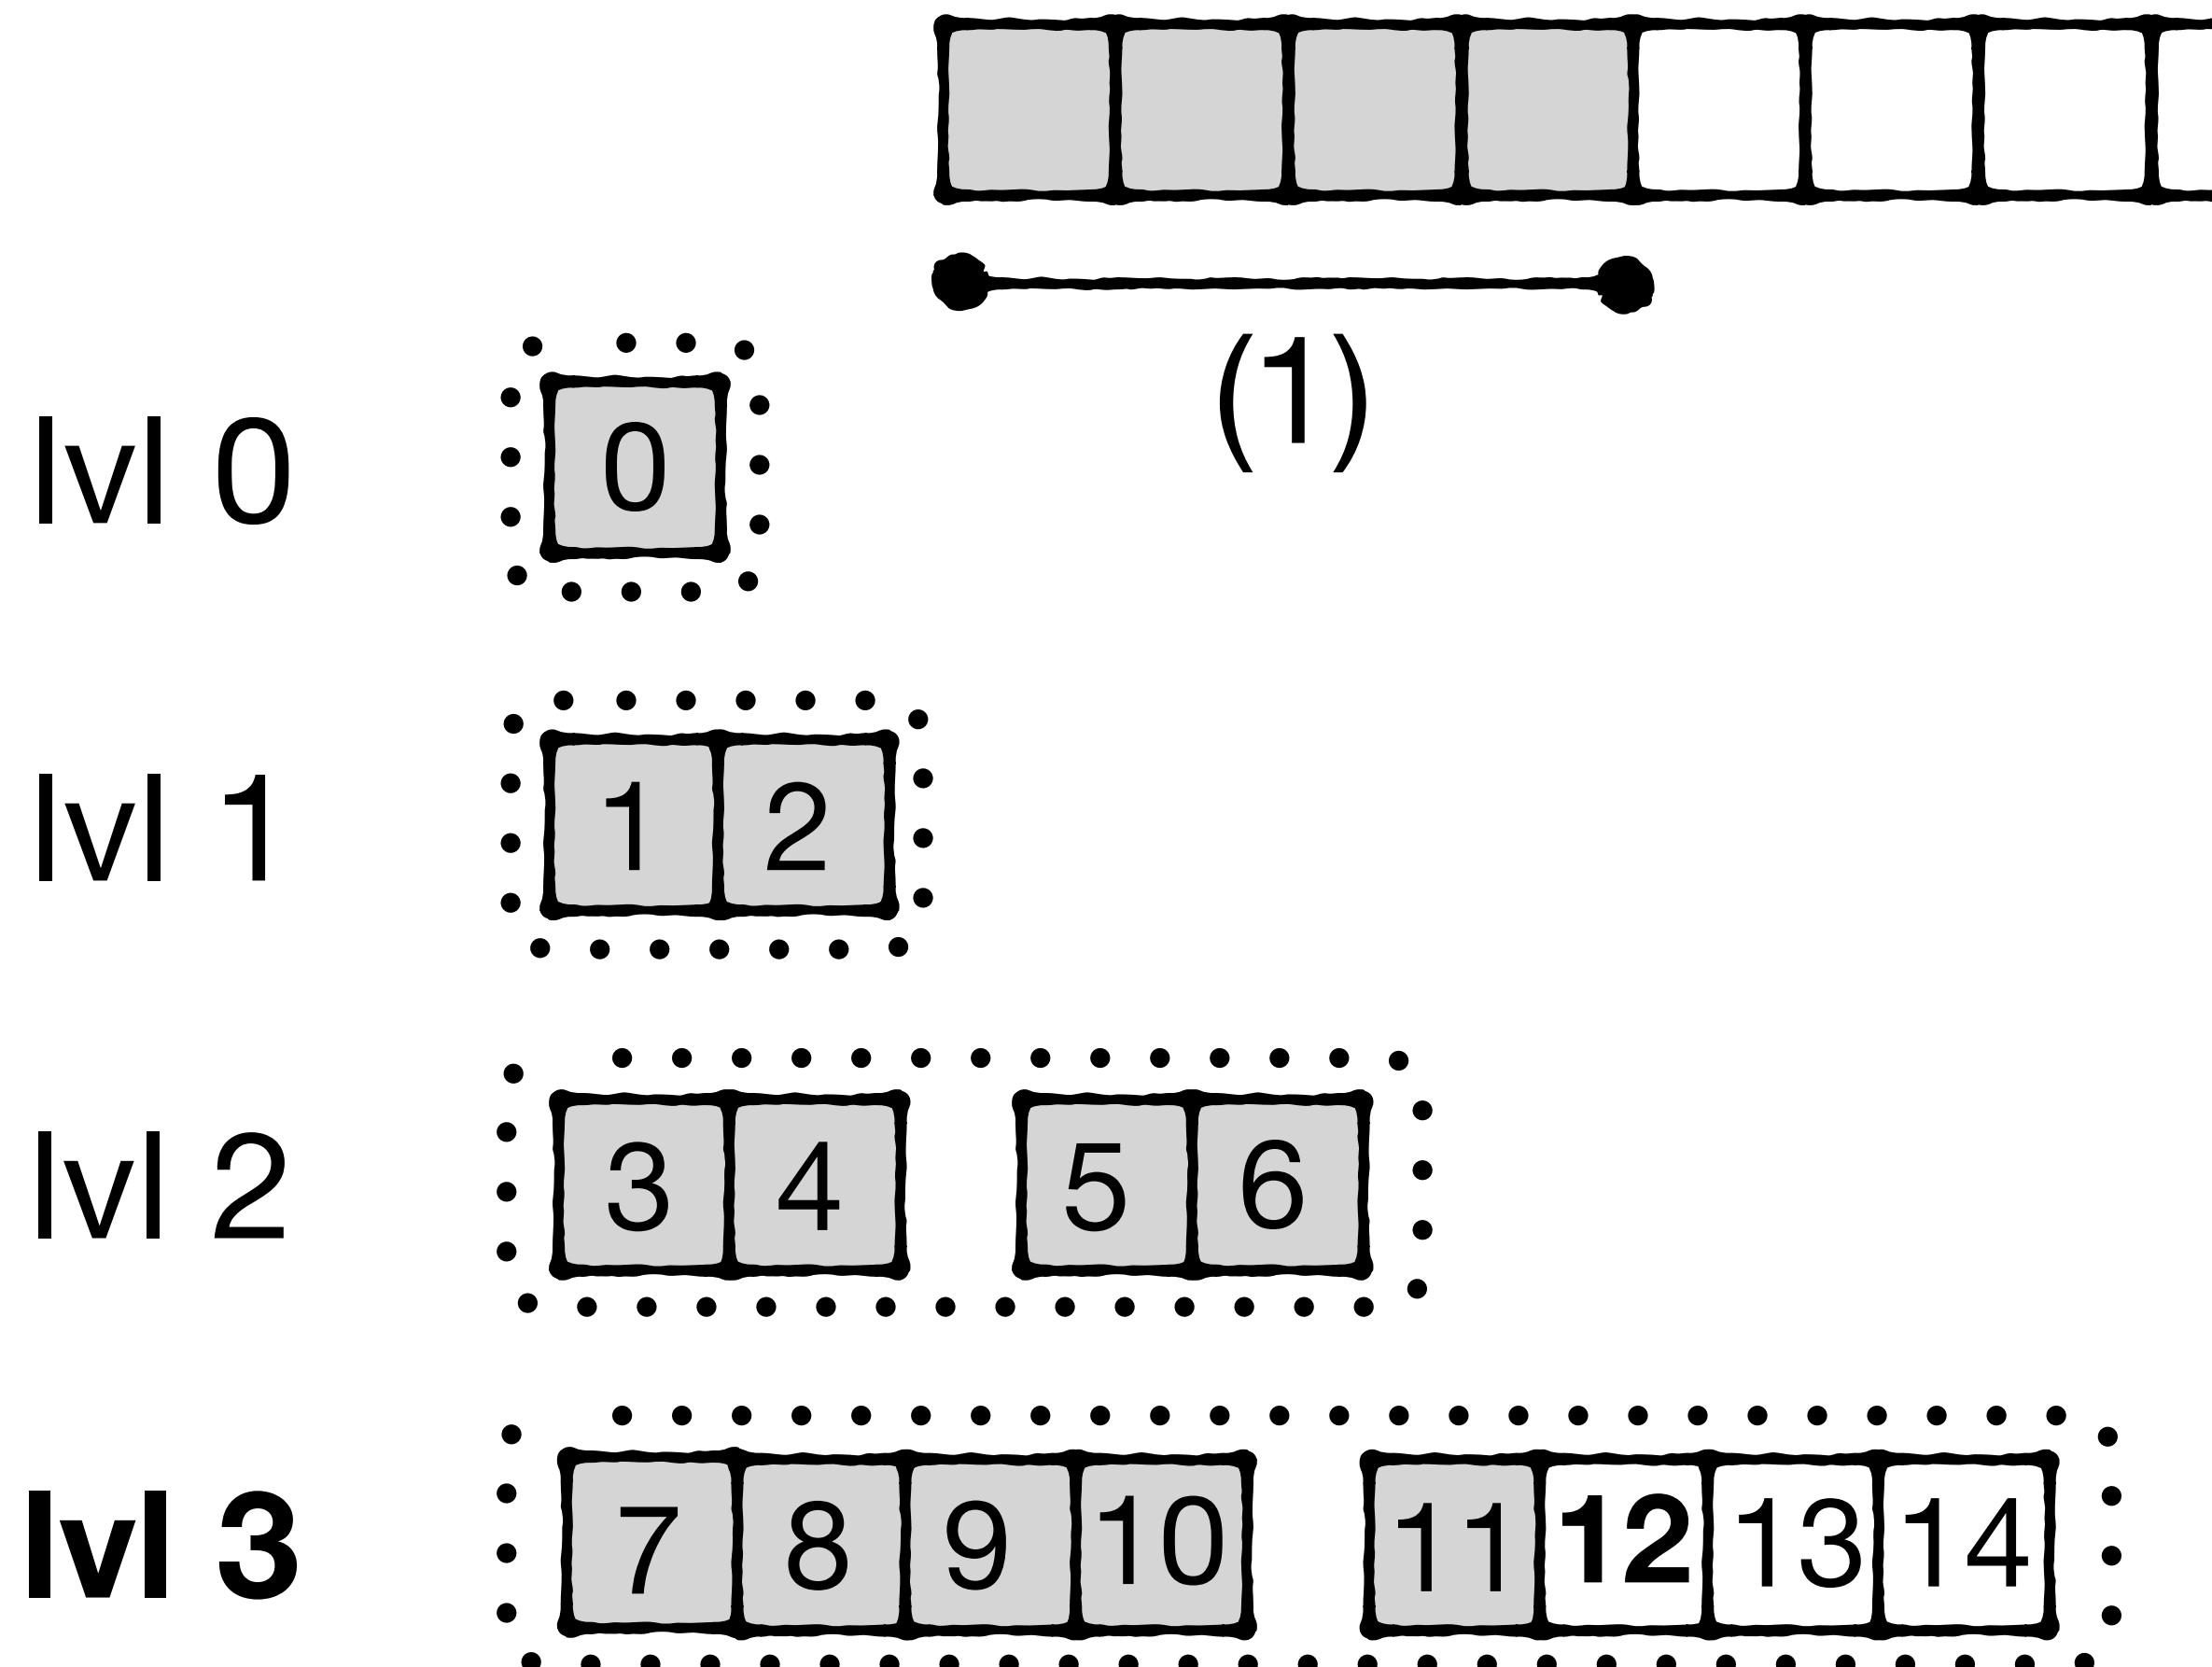
$$= 2^{\lfloor k/2 \rfloor} \cdot (2 + (k \bmod 2)) - 2$$

$$= (1 \ll (k \gg 1)) \cdot (2 + (k \& 1)) - 2$$



Shift to left

get(i)



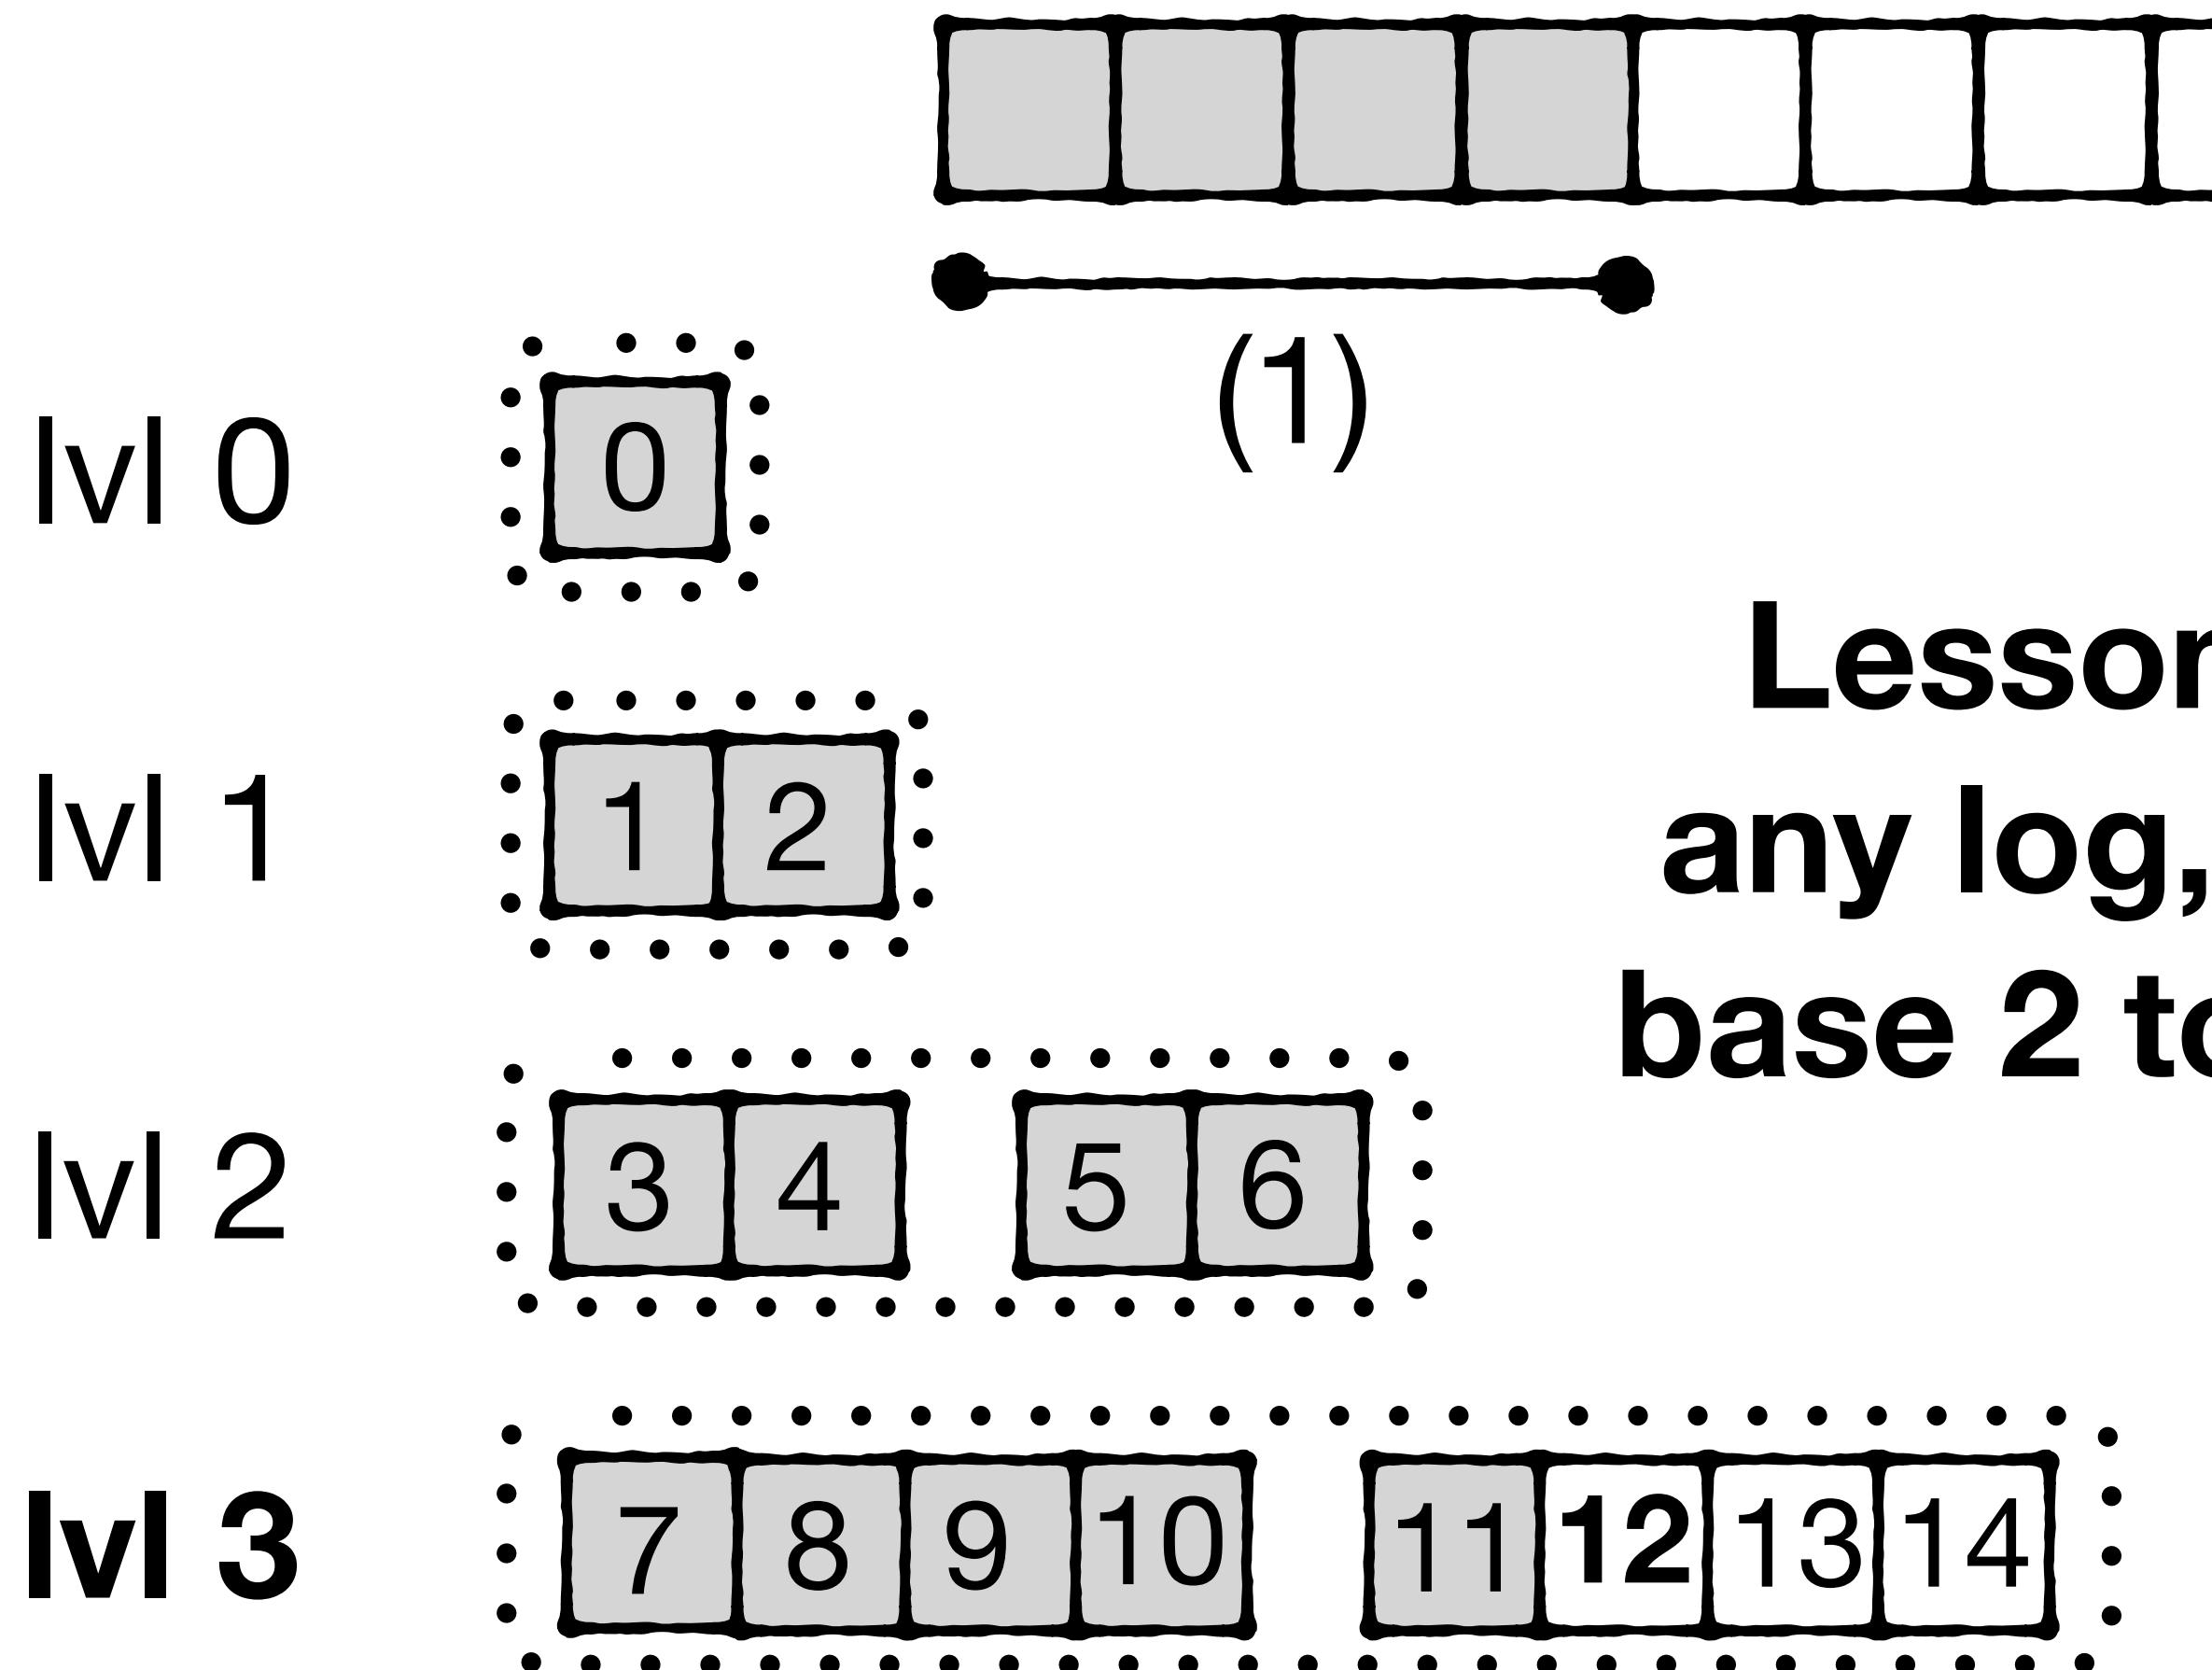
blocks in levels 0 to k-1

$$= 2^{\lfloor k/2 \rfloor} \cdot (2 + (k \bmod 2)) - 2$$

$$= (1 \ll (k \gg 1)) \cdot (2 + (k \& 1)) - 2$$

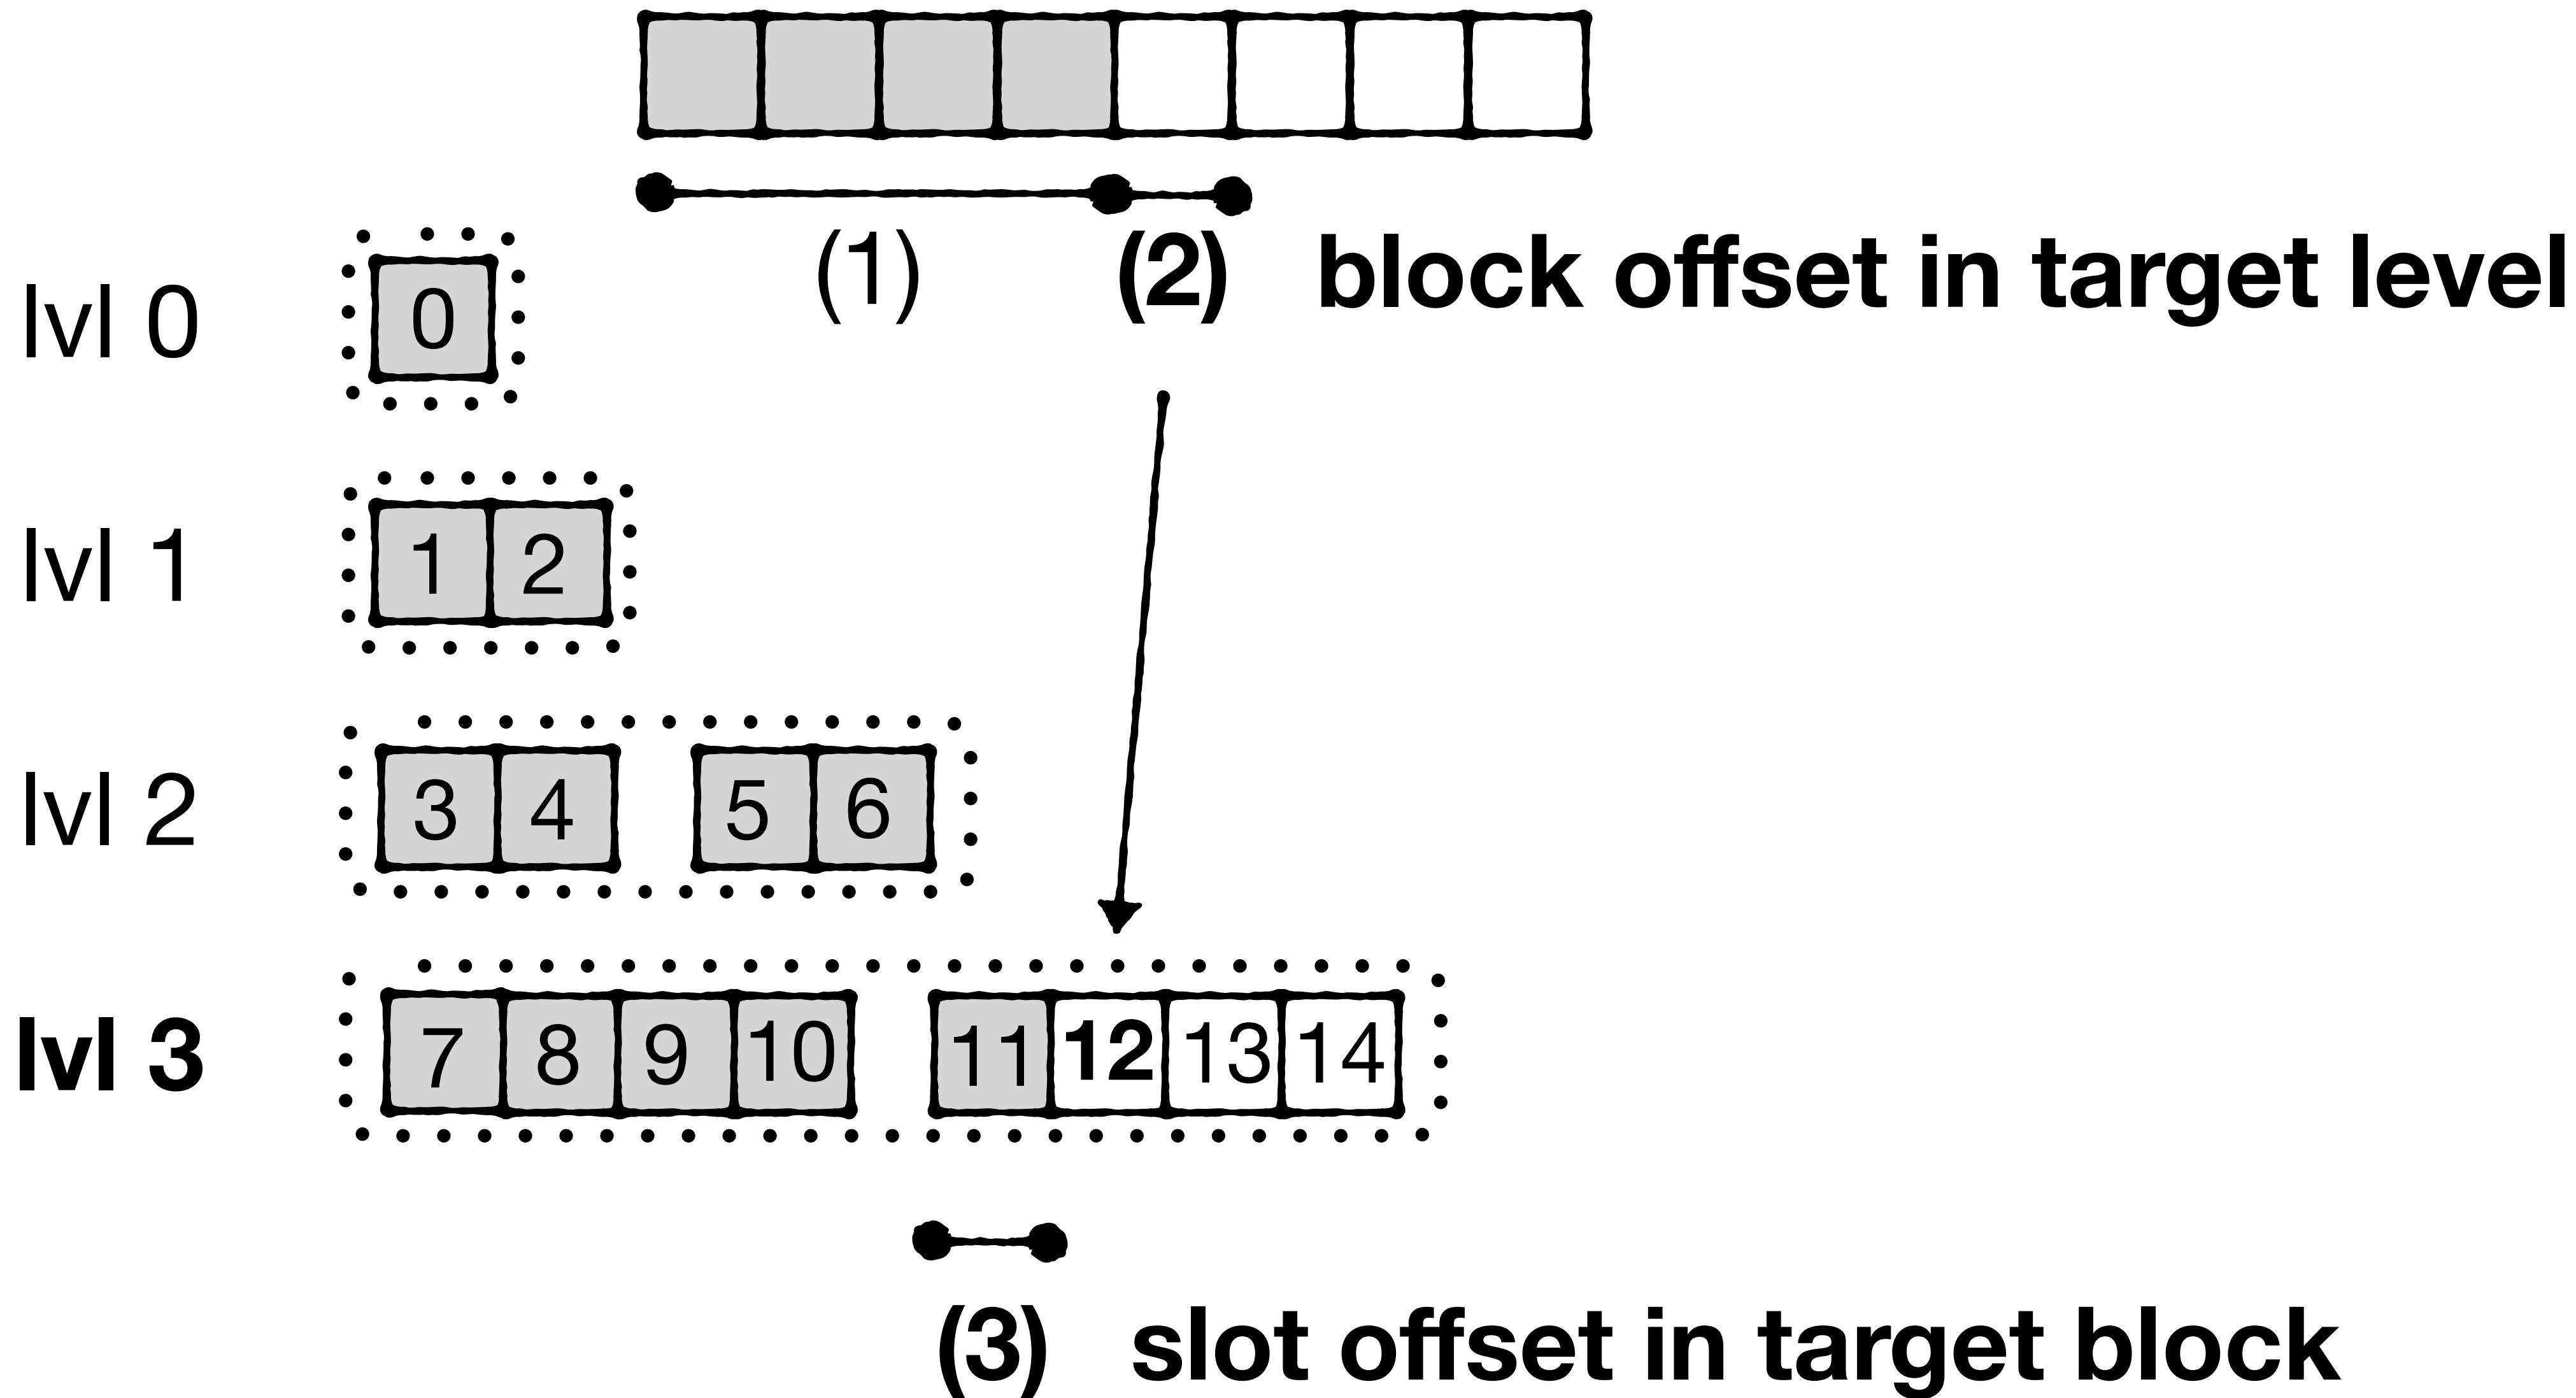
≈0.6 ns rather than ≈10ns

get(i)

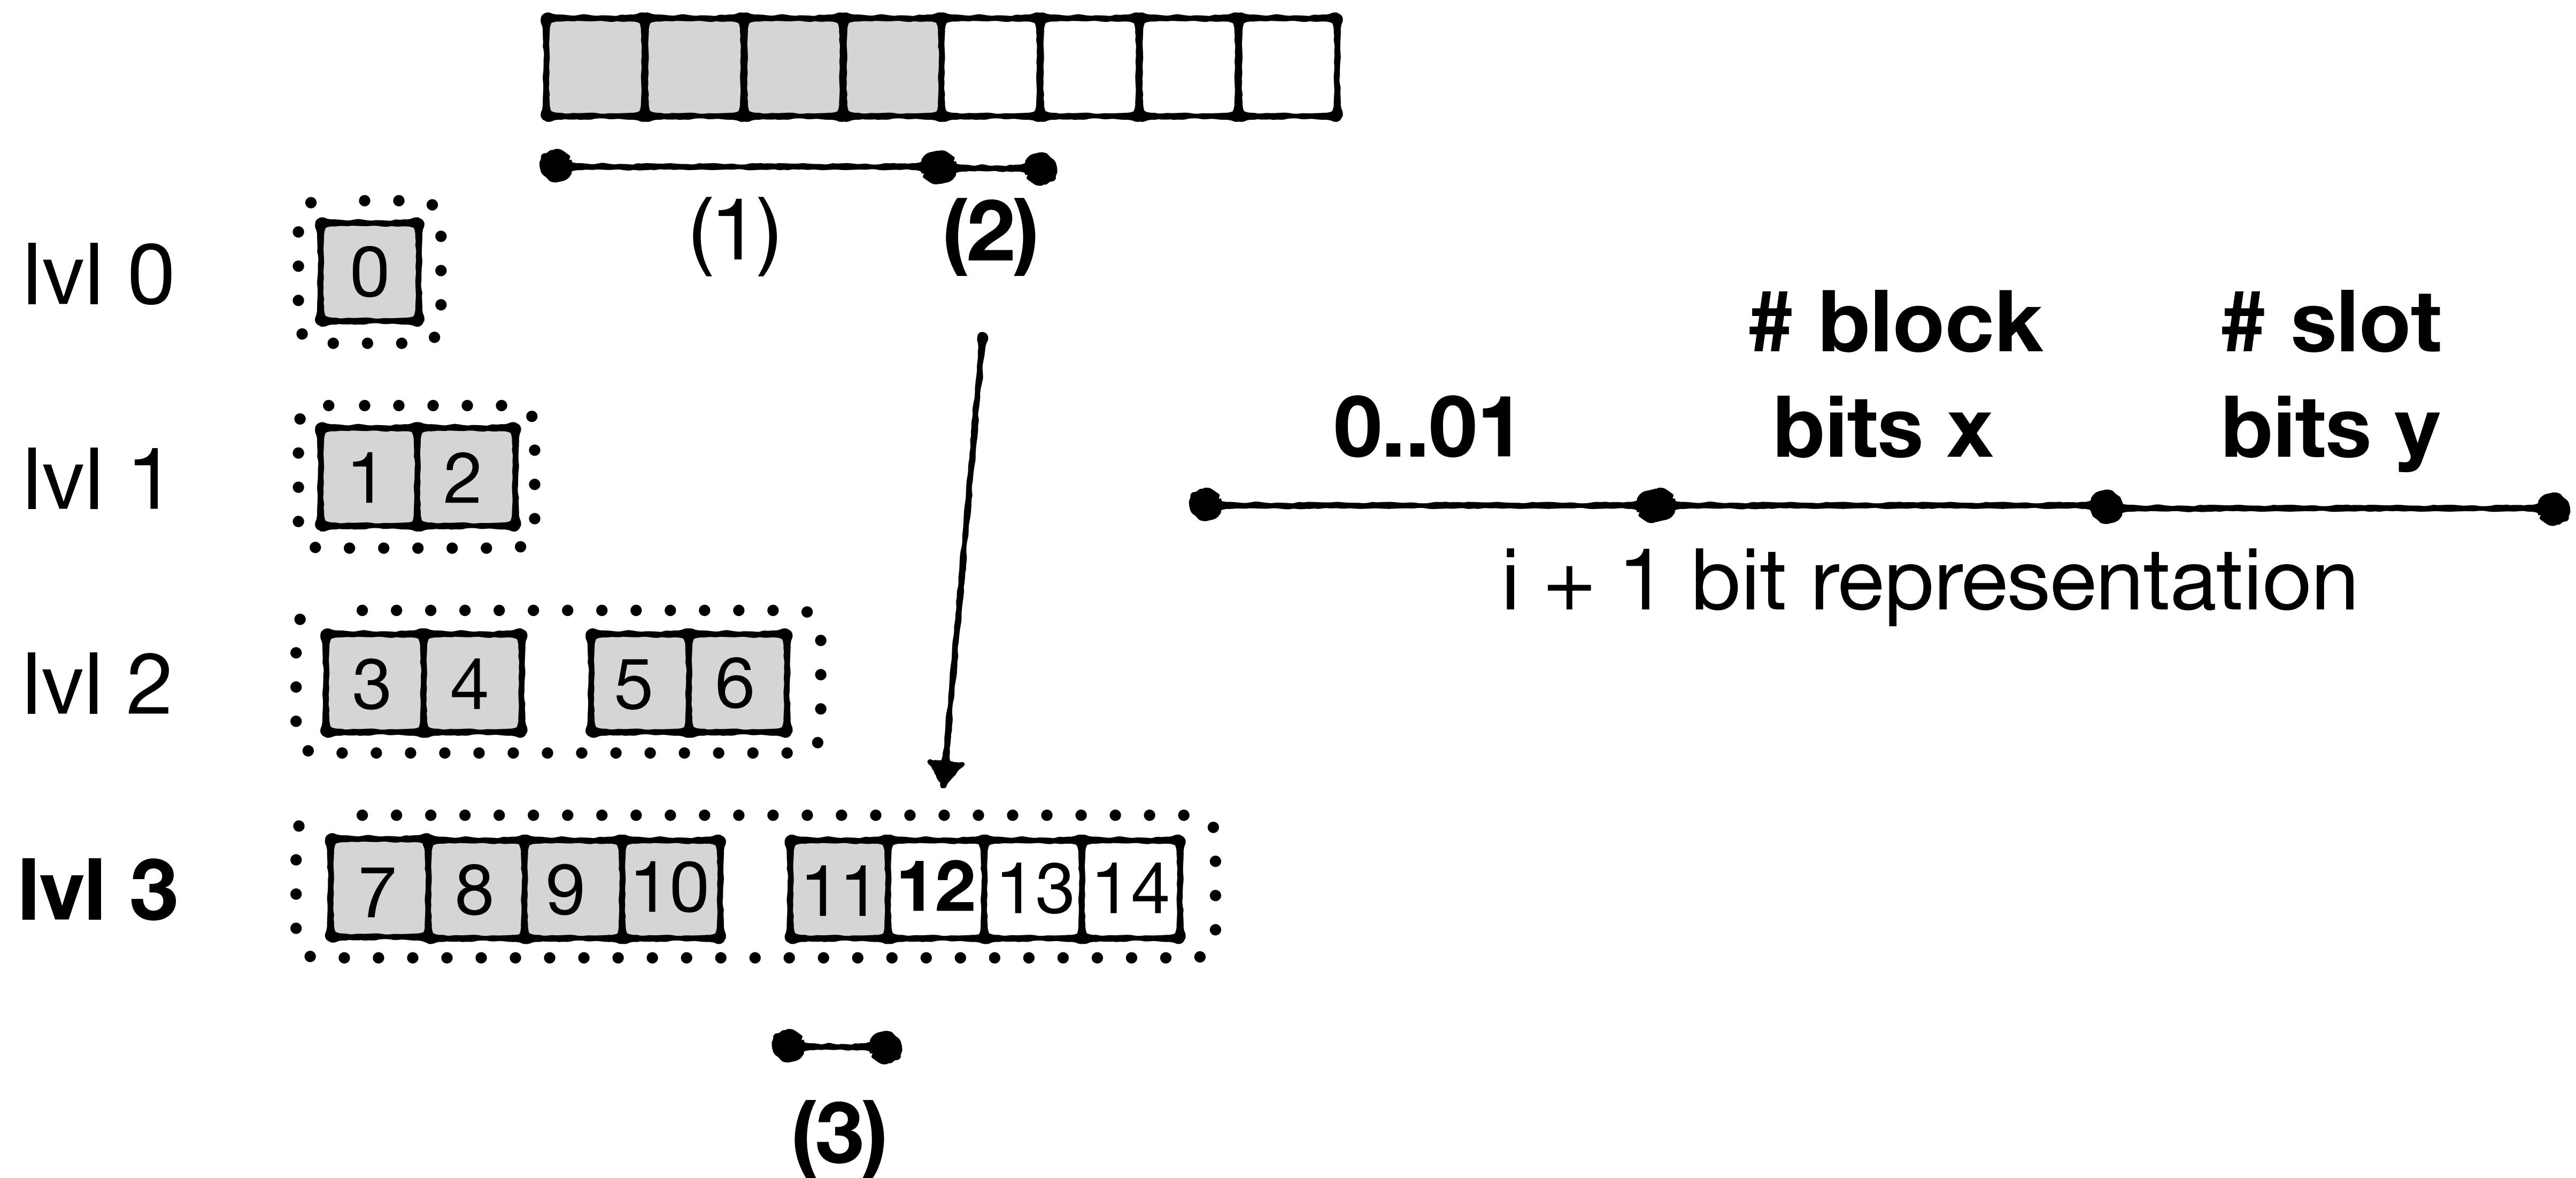


**Lesson: design structure such that
any log, division, or exponentiation is
base 2 to support fast CPU operations**

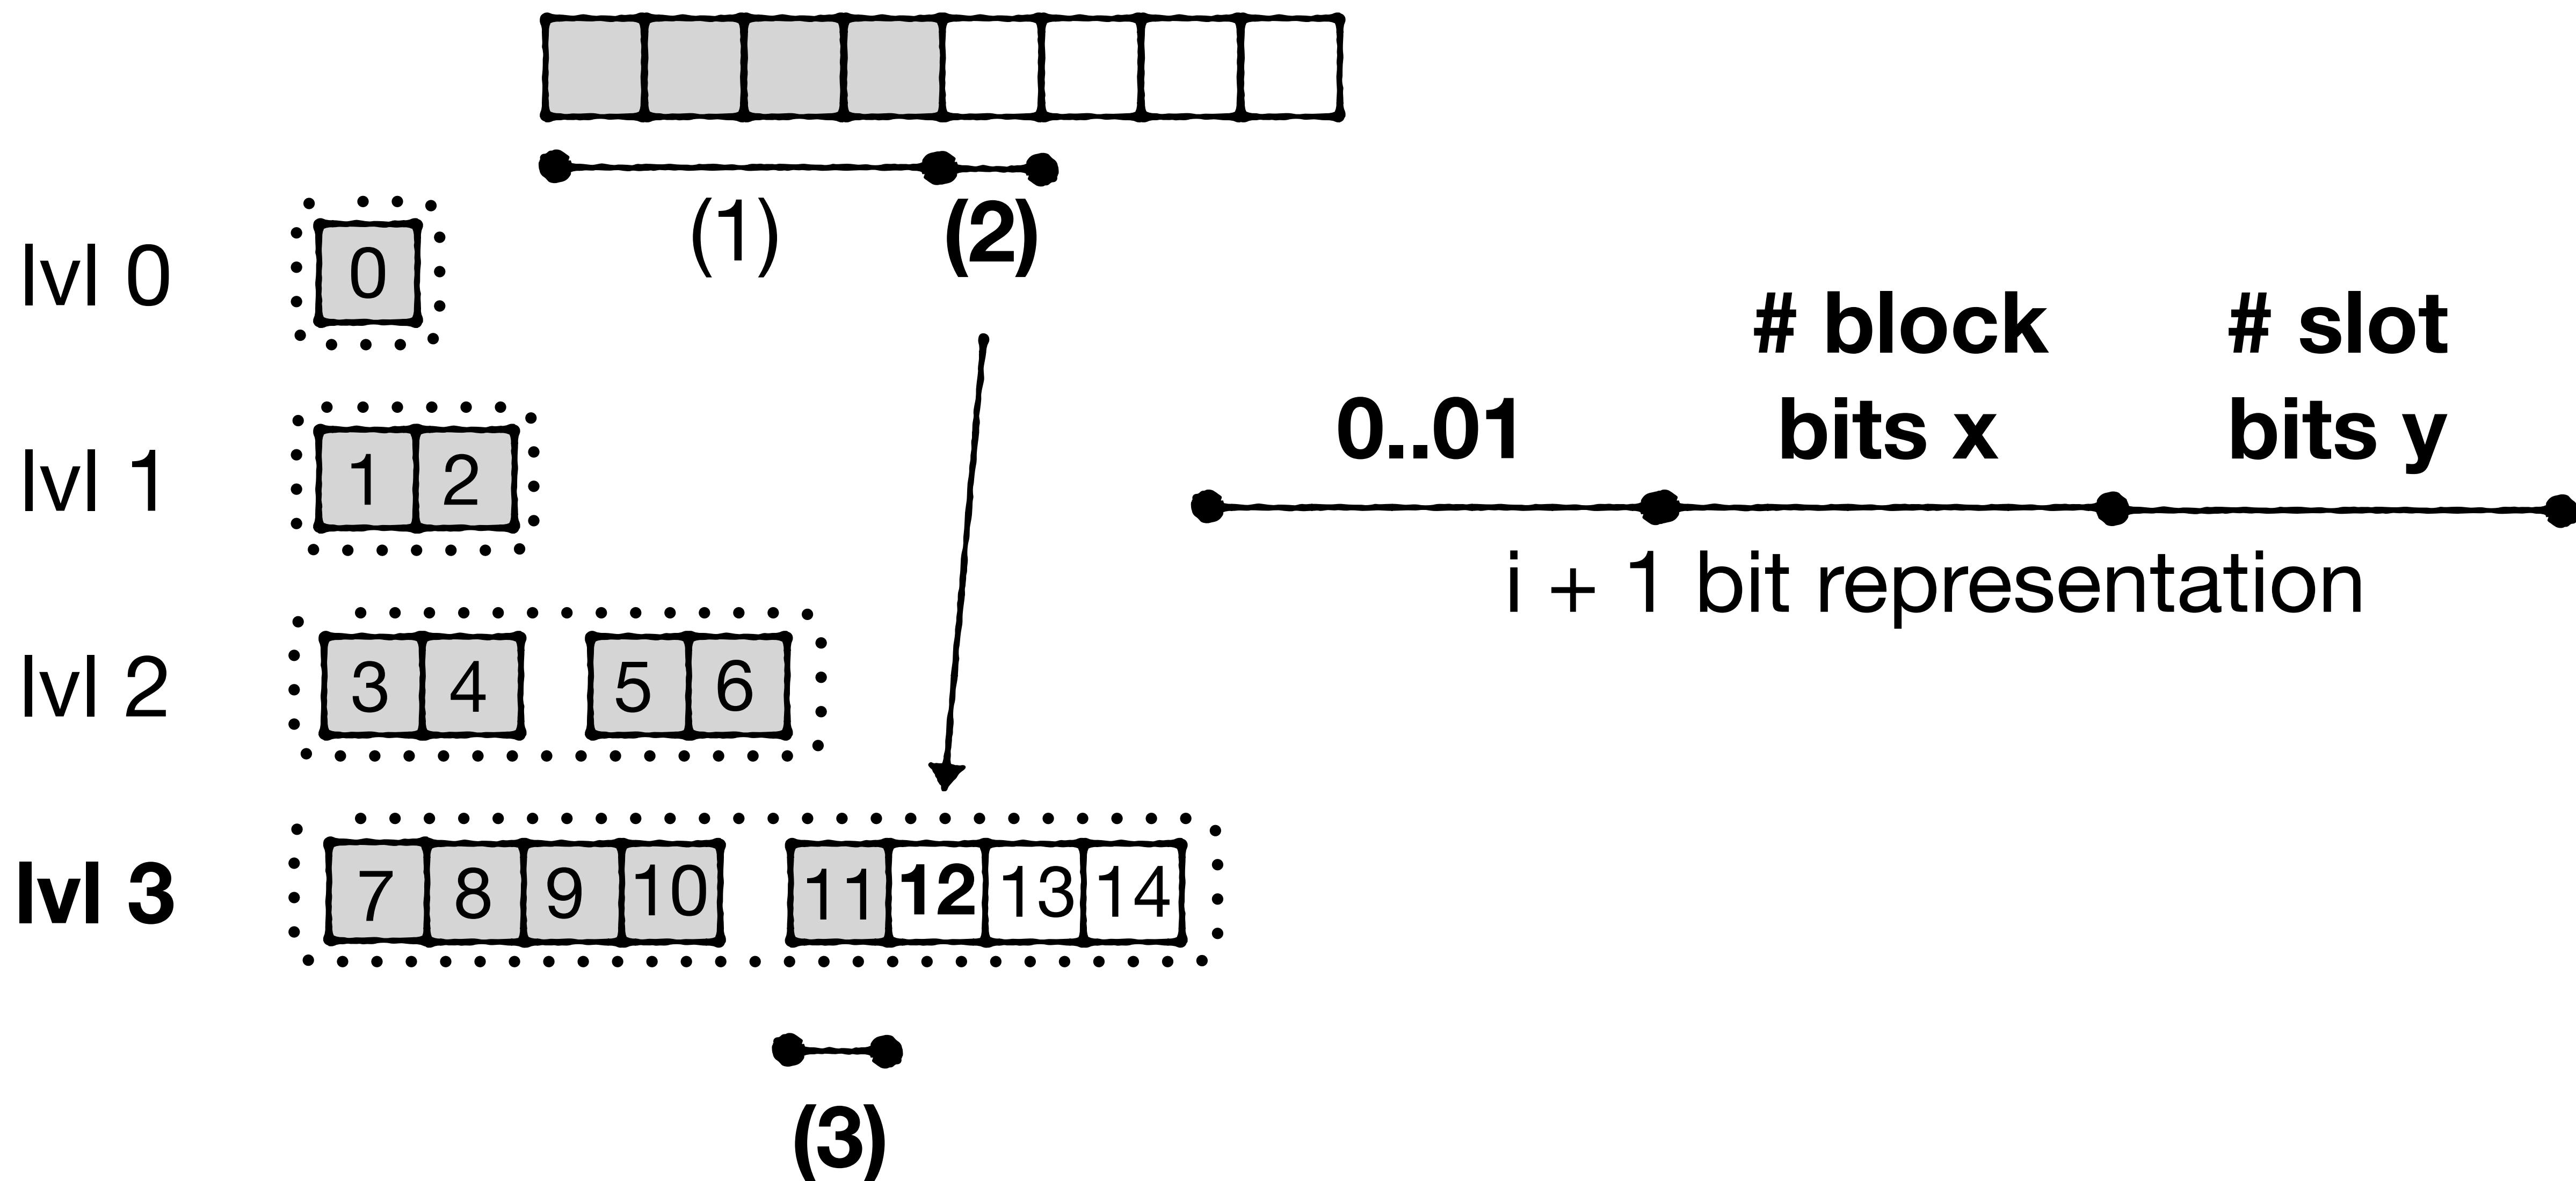
get(i)



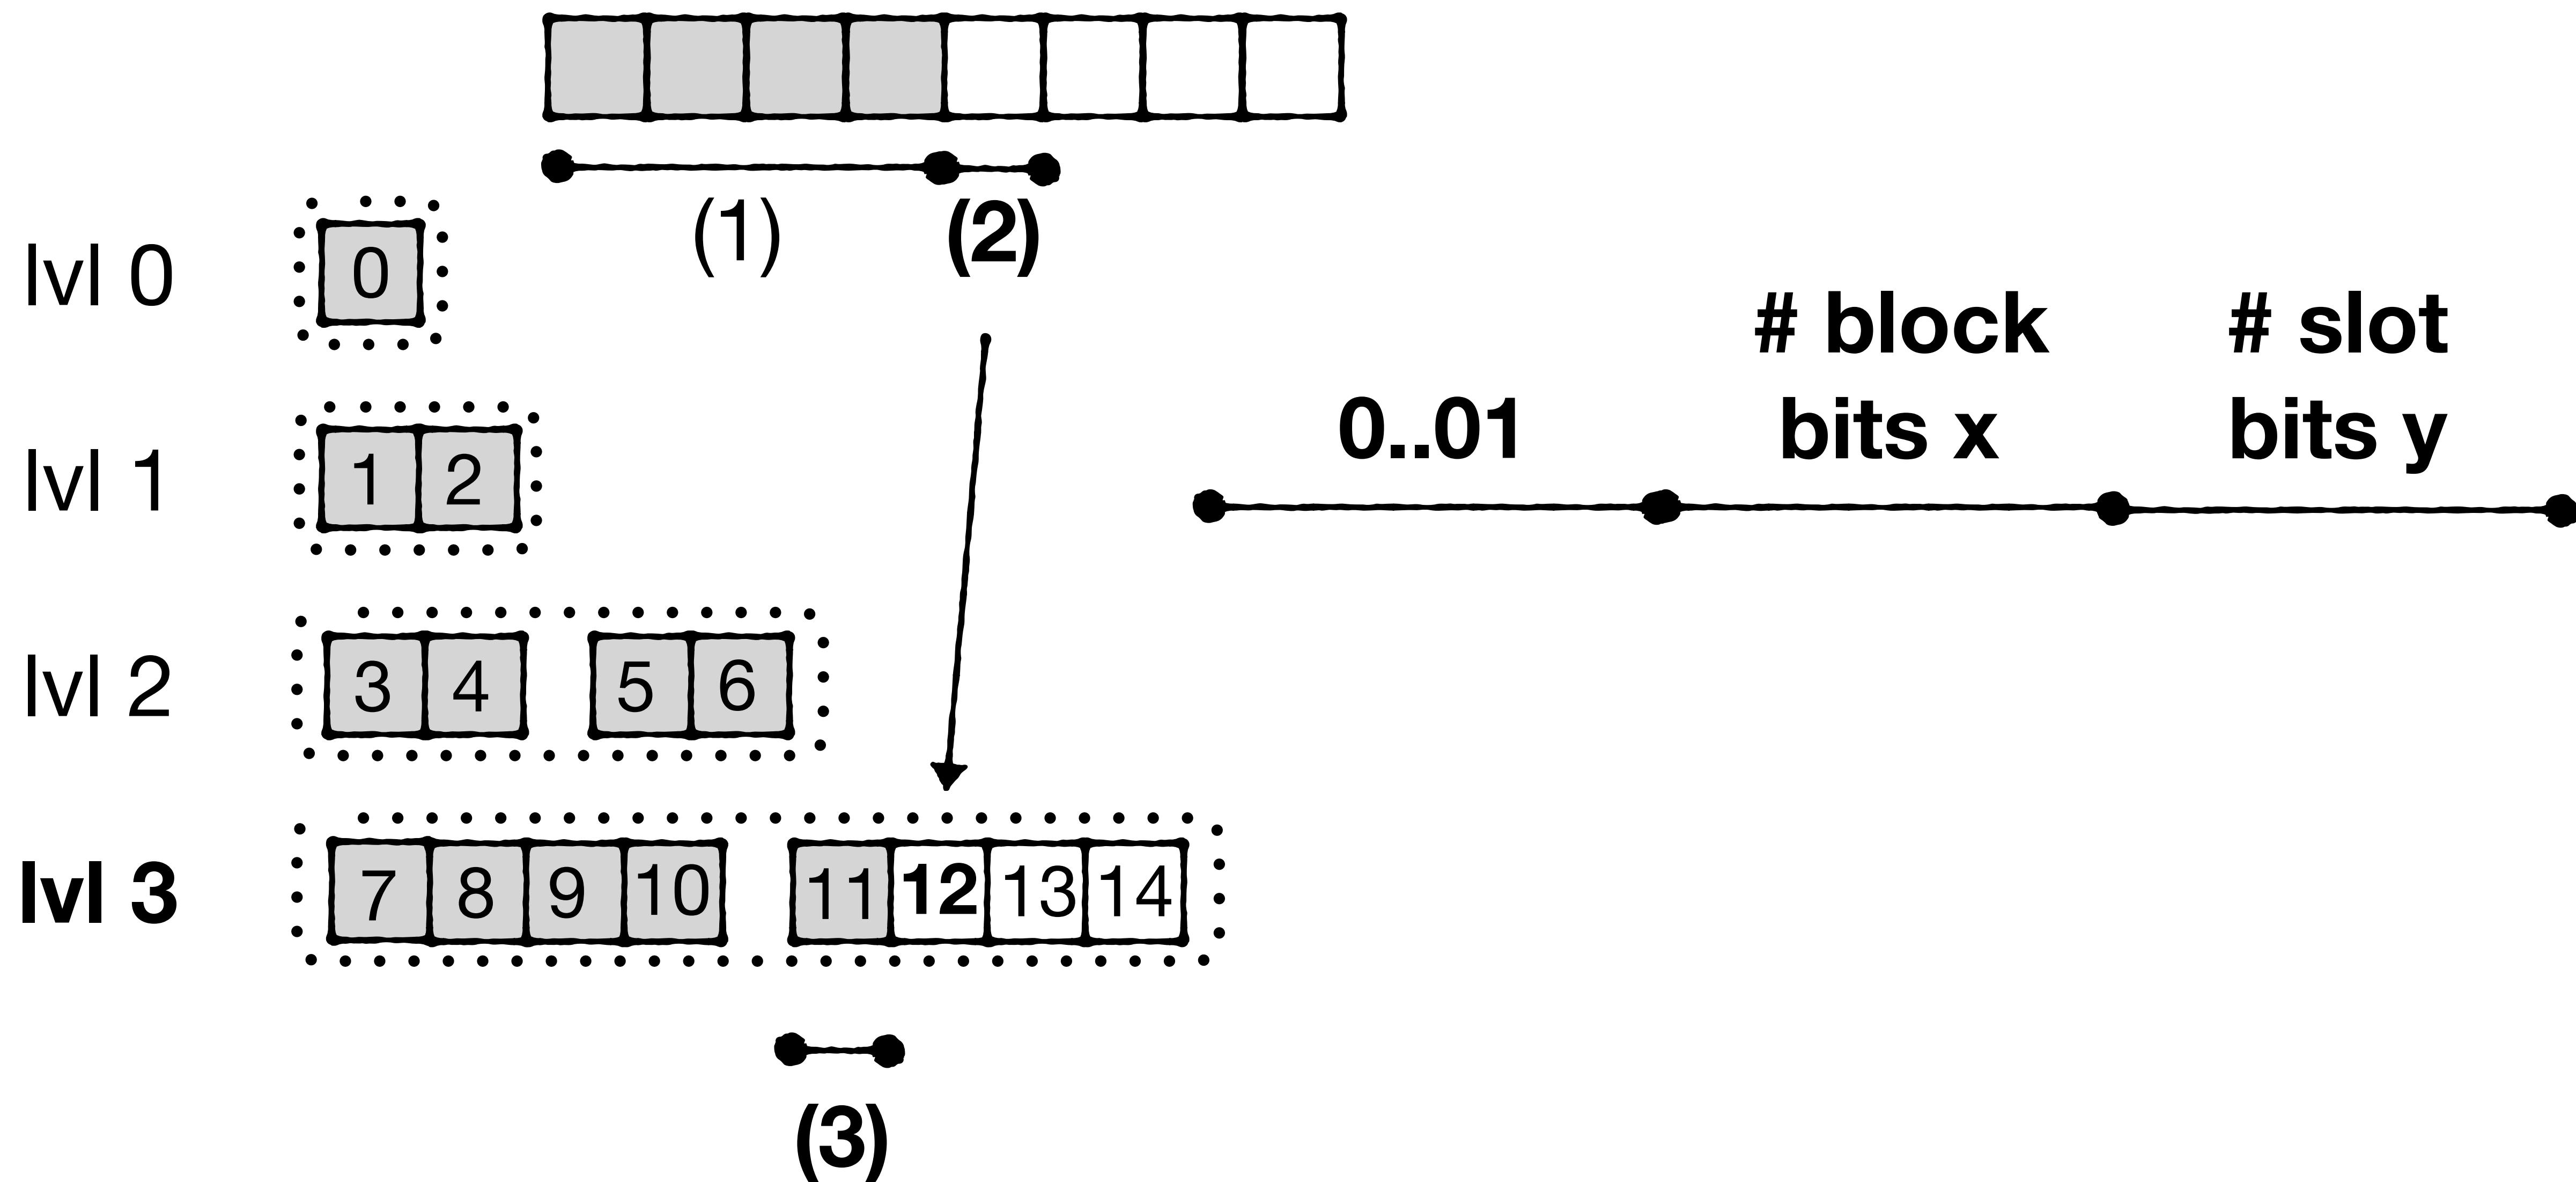
get(*i*)



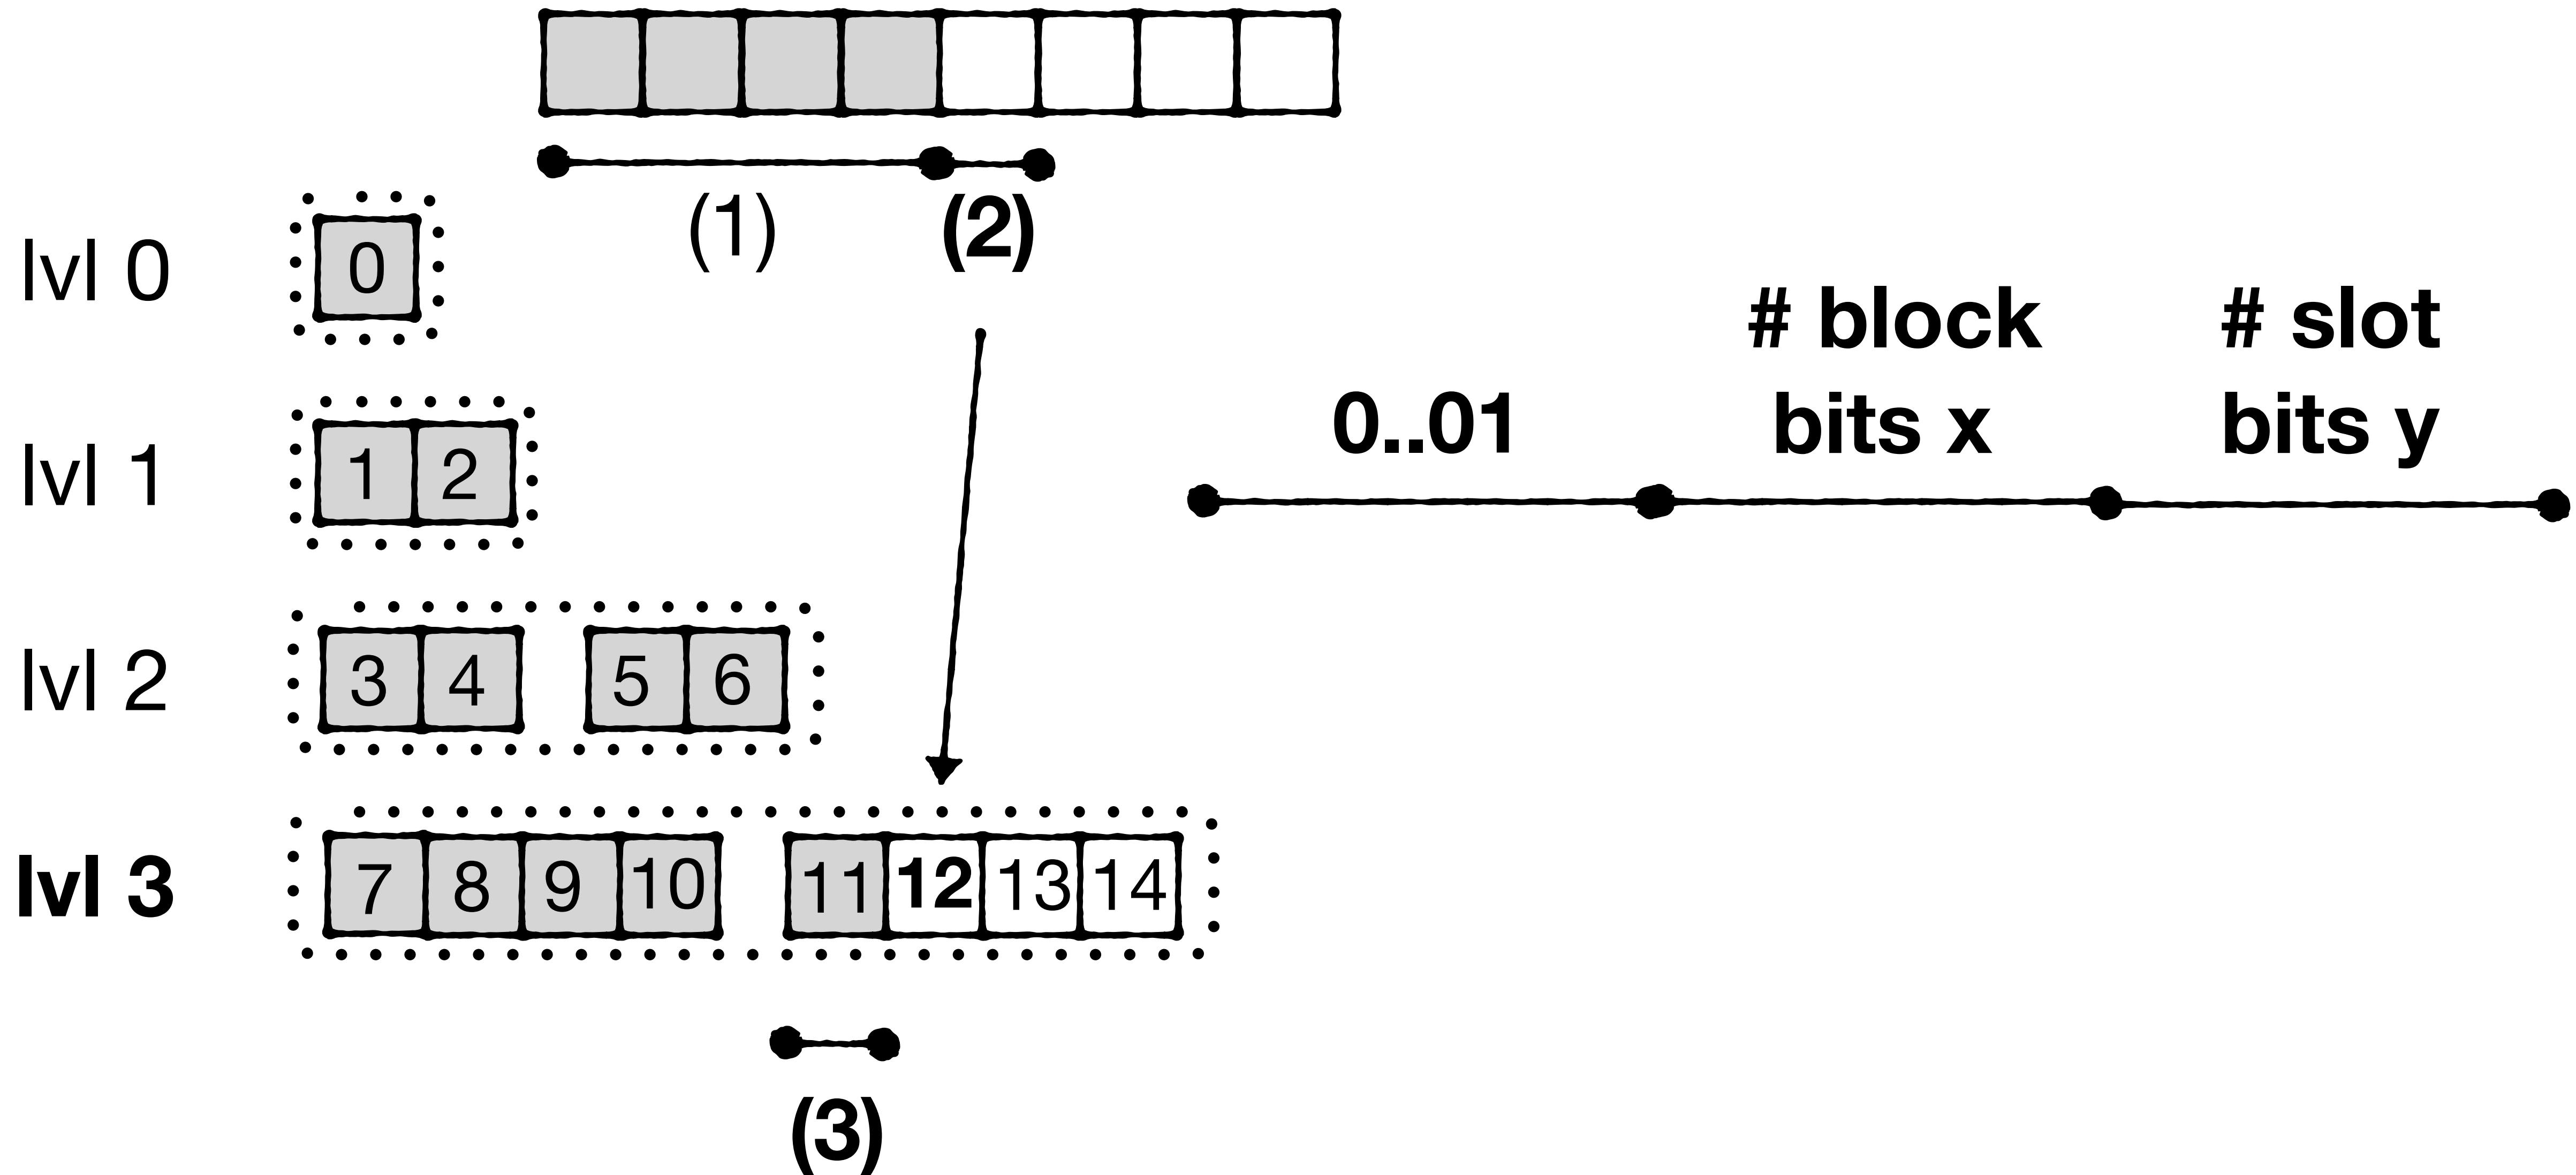
get(00001100)

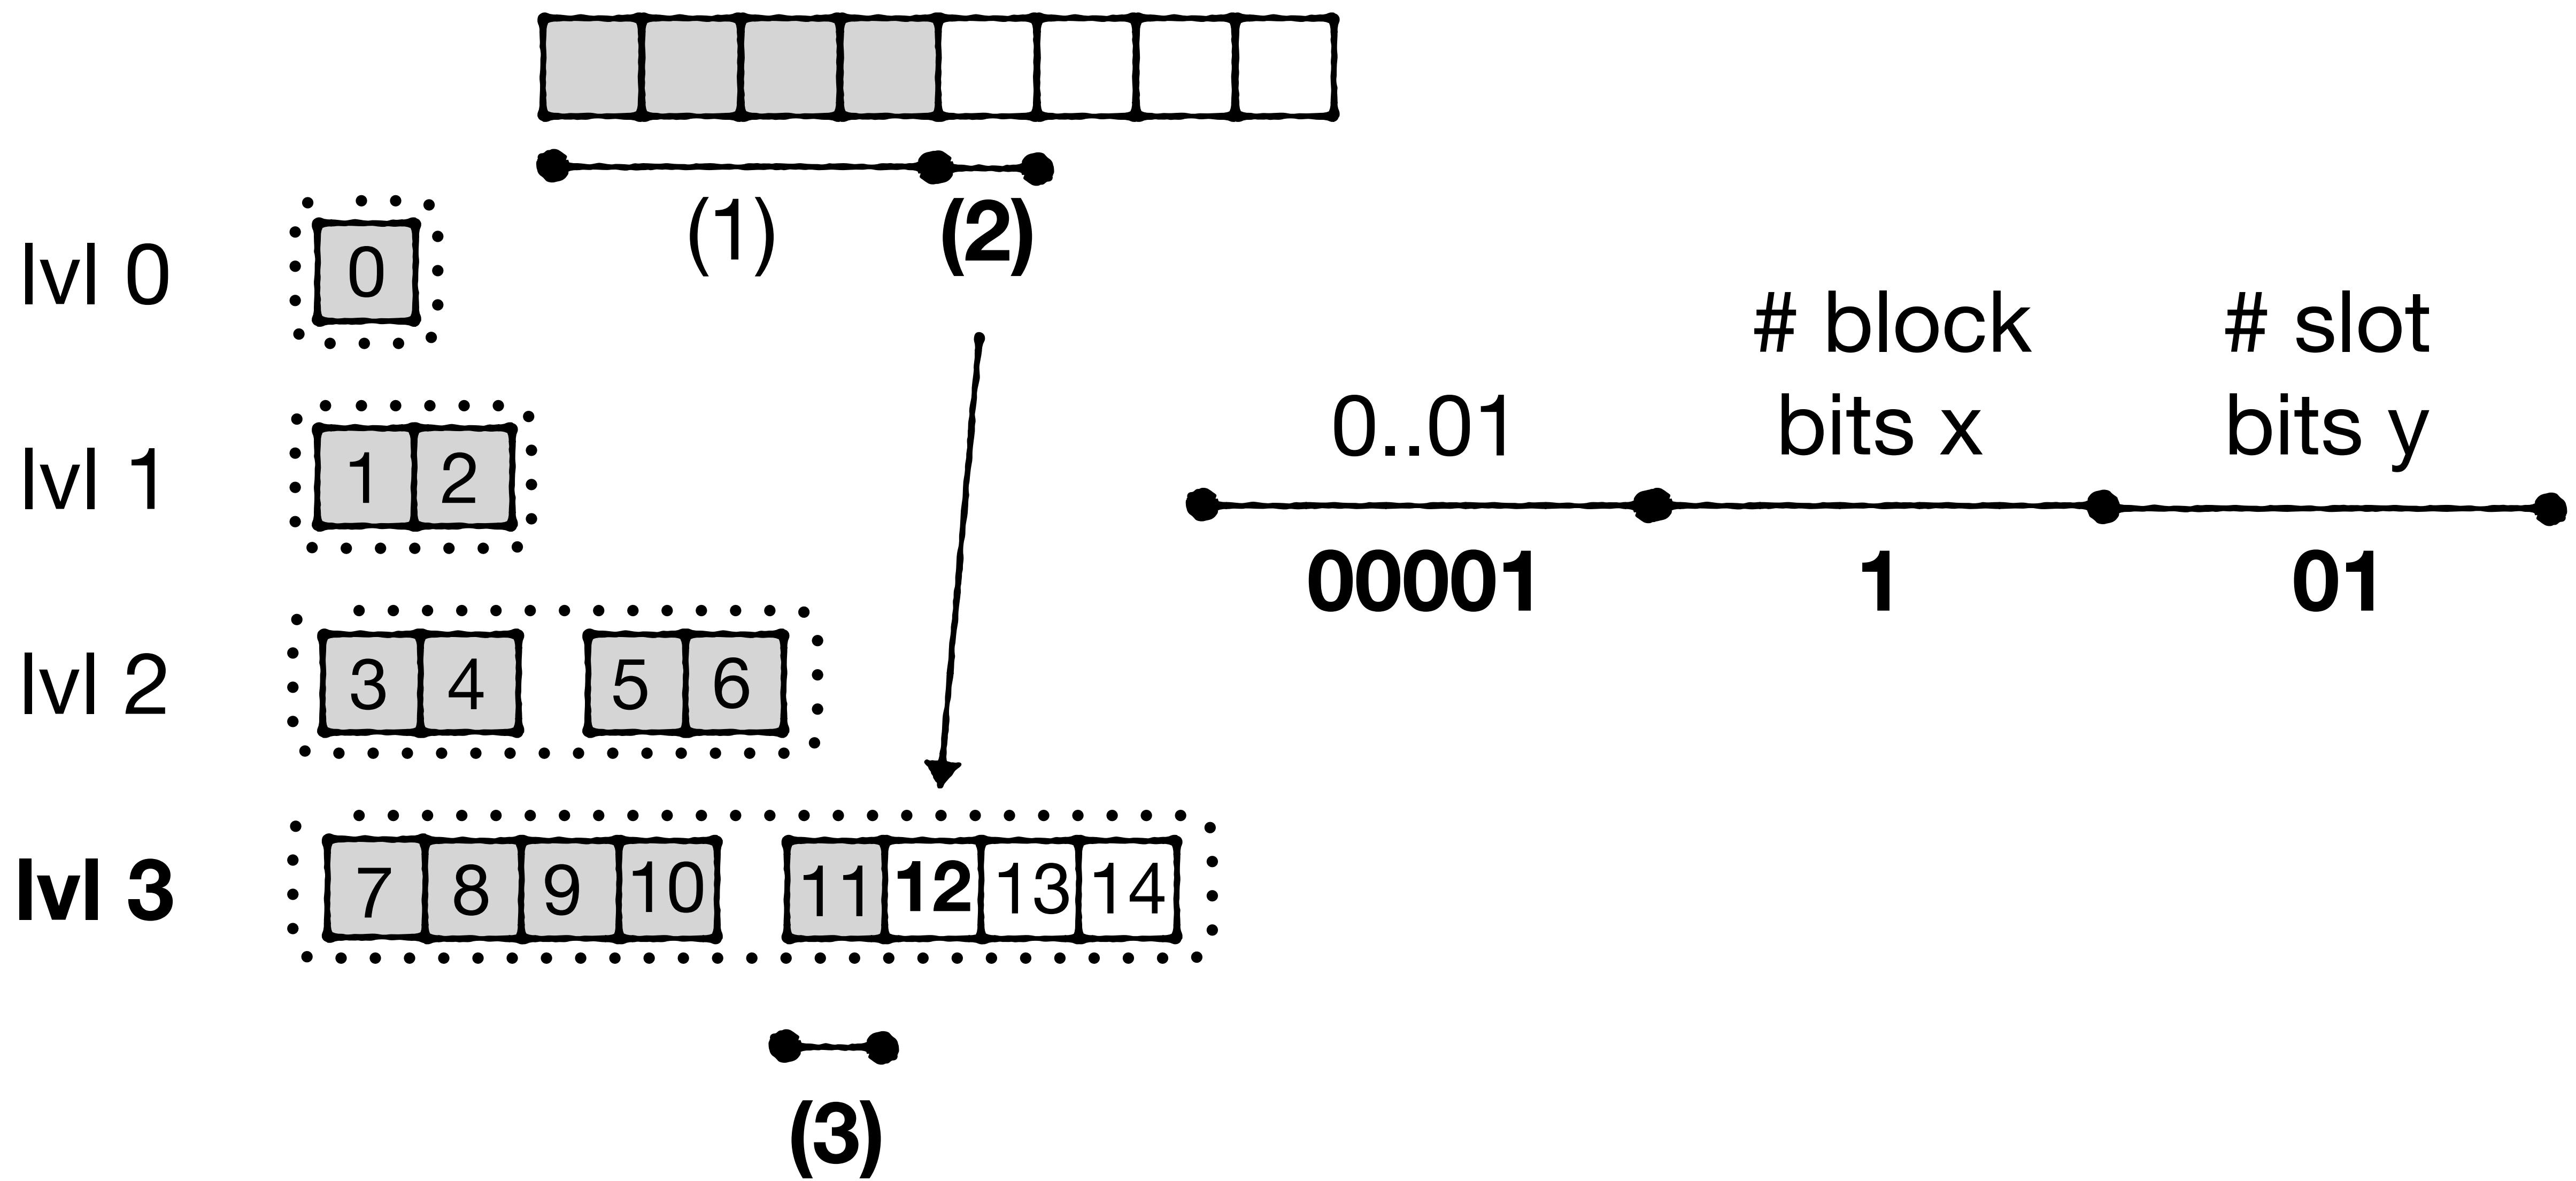


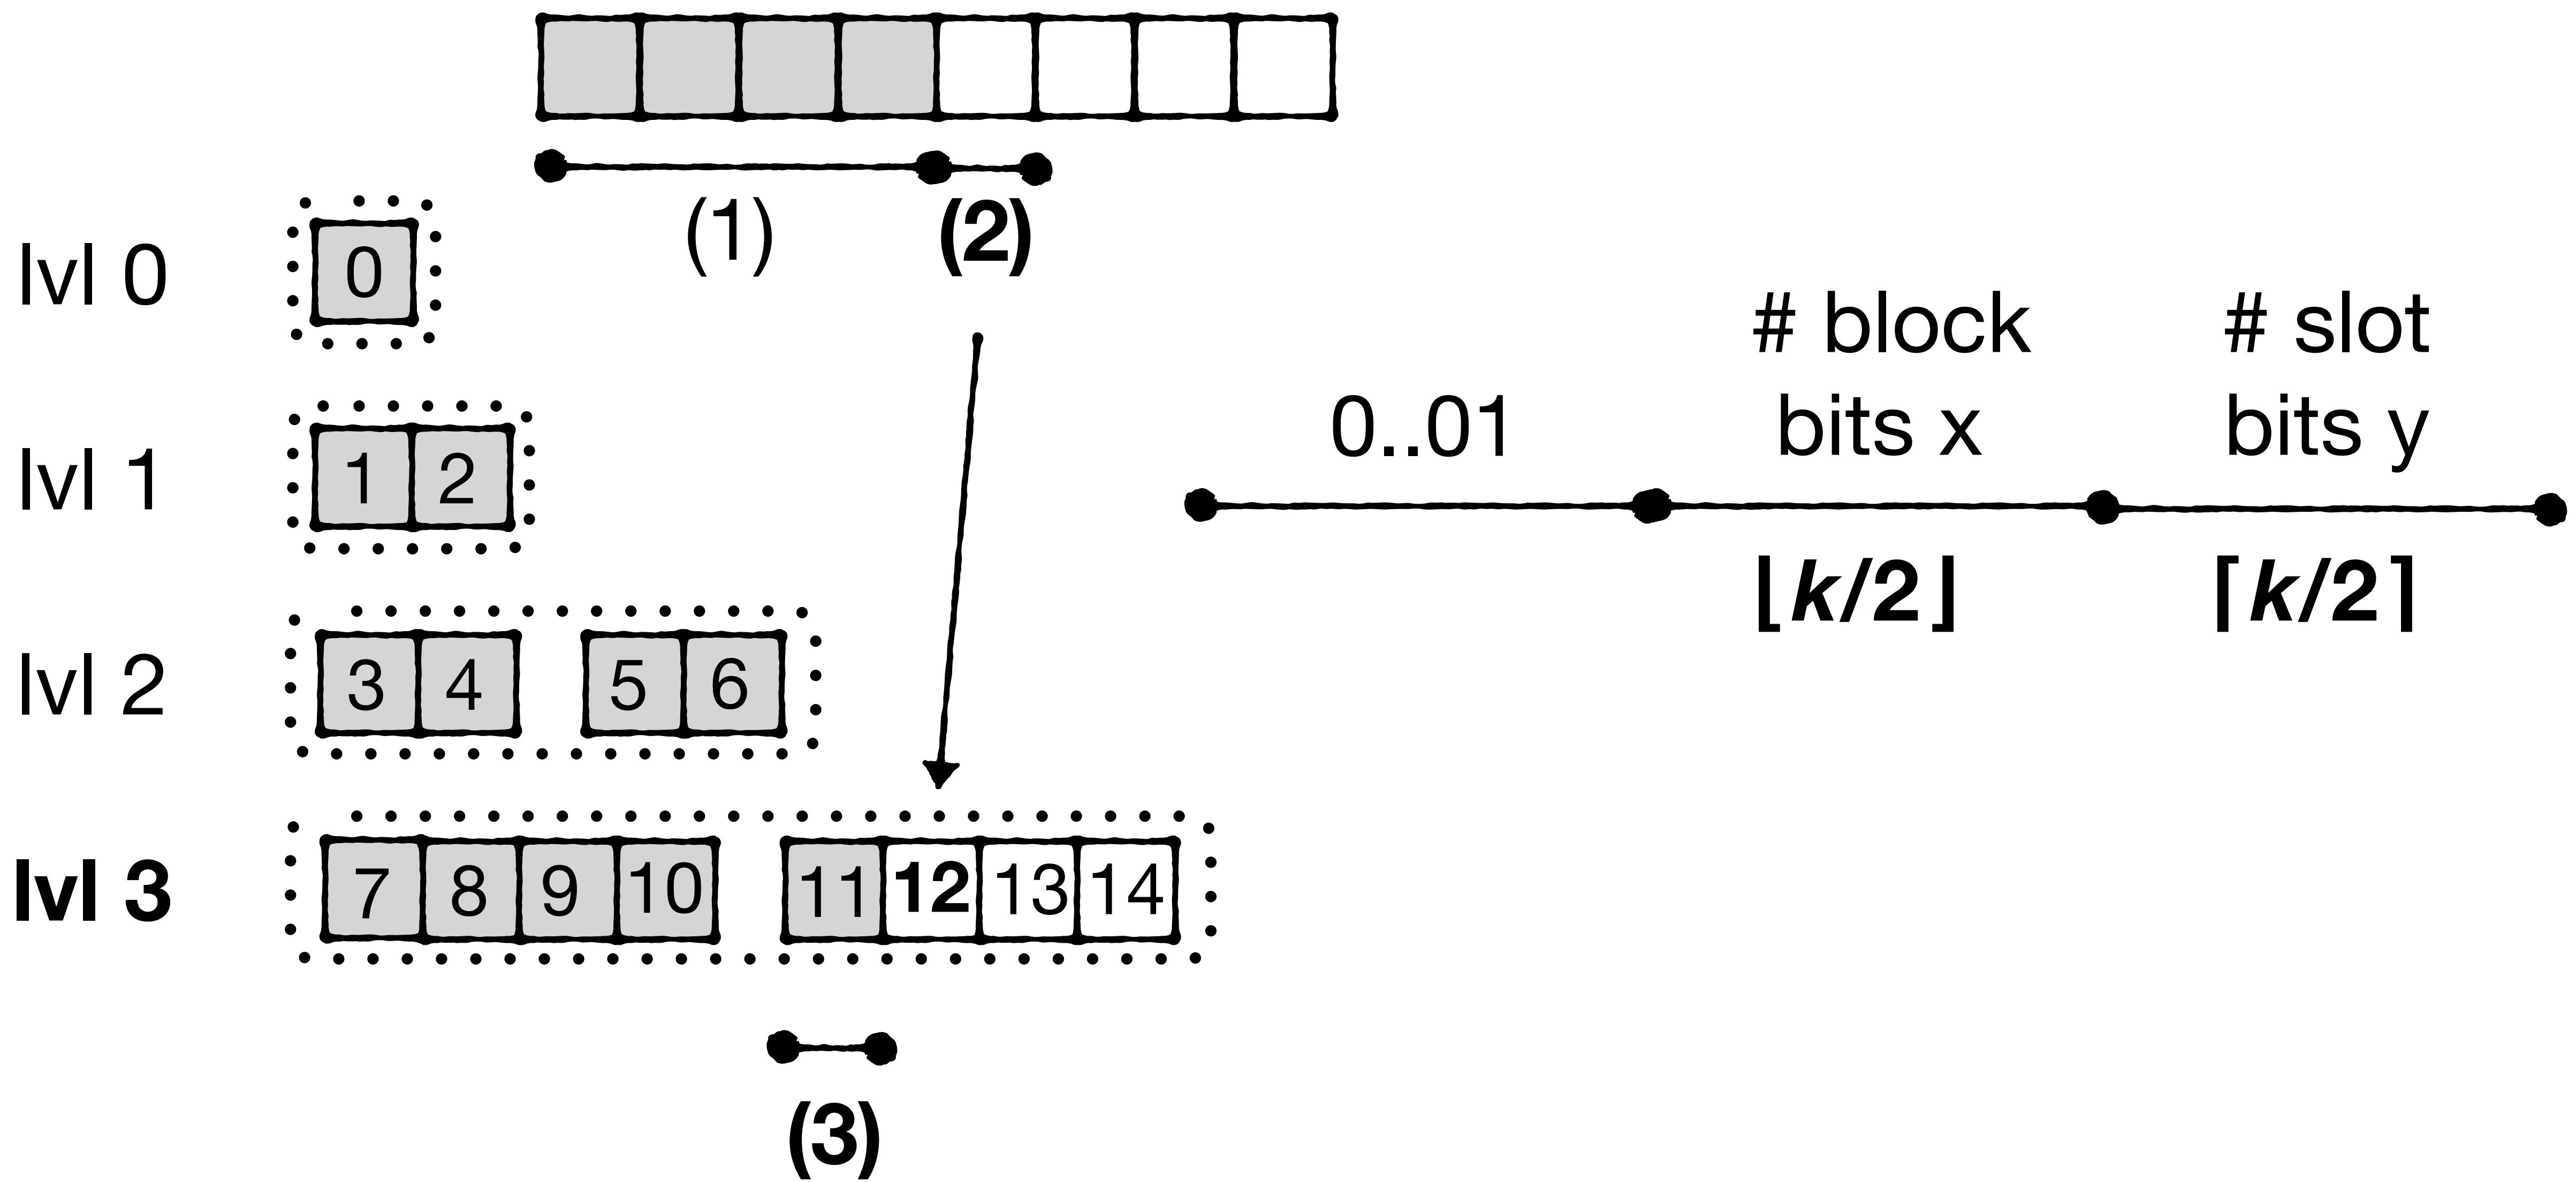
00001100 + 1

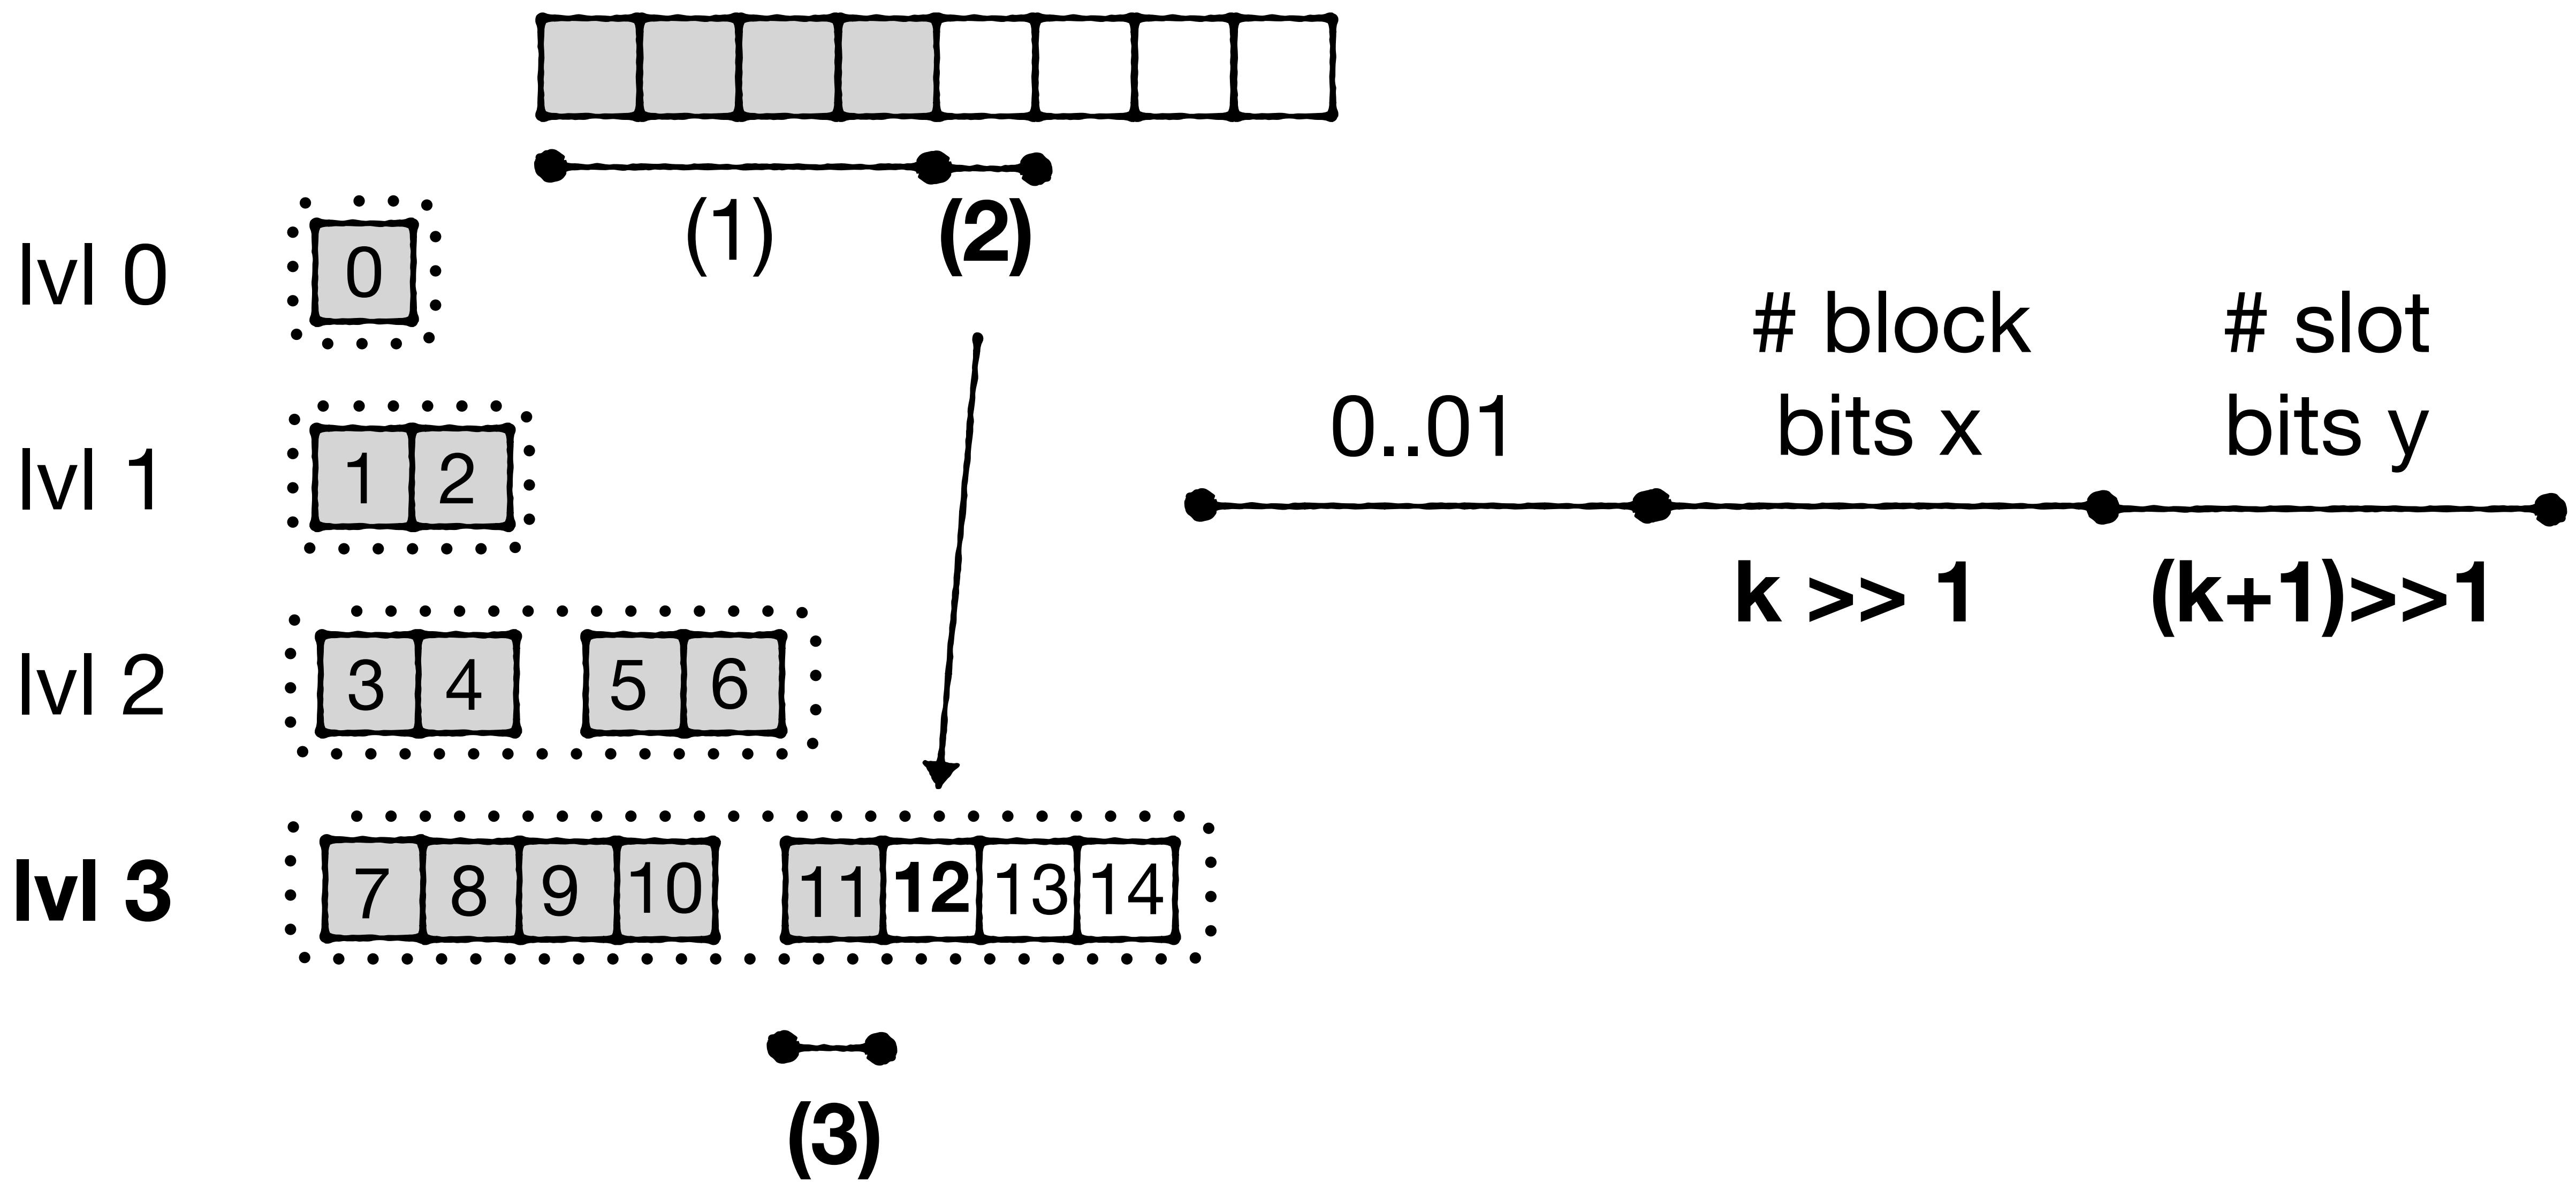


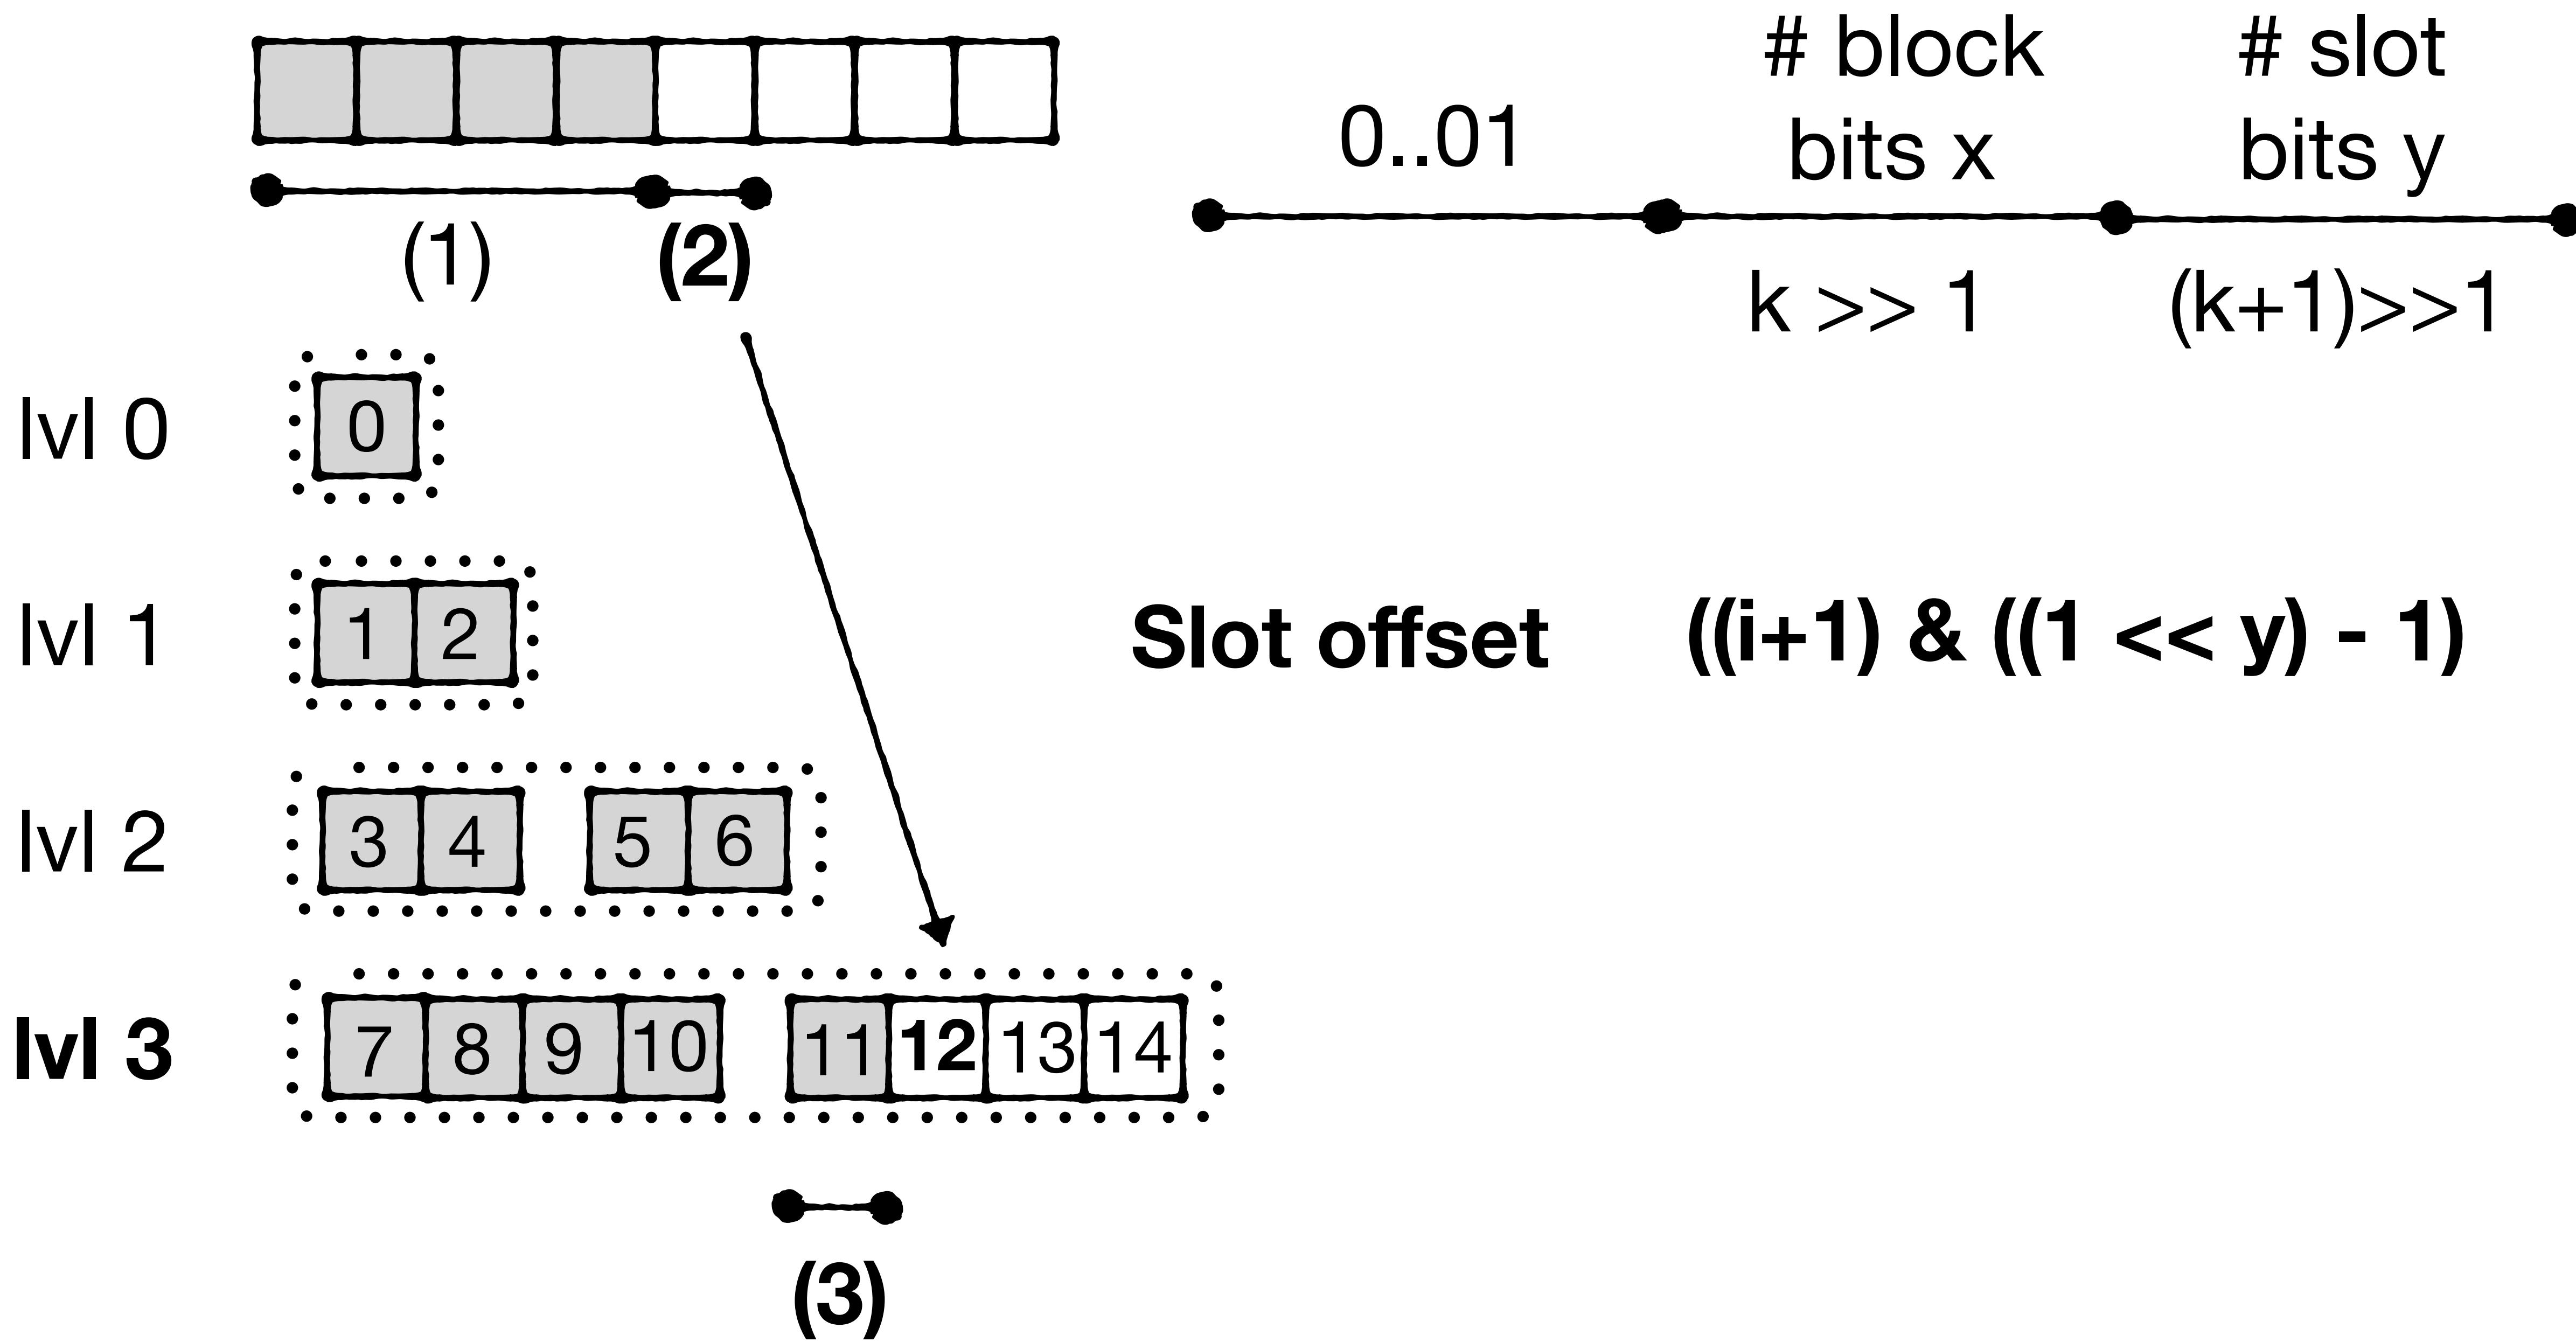
00001101

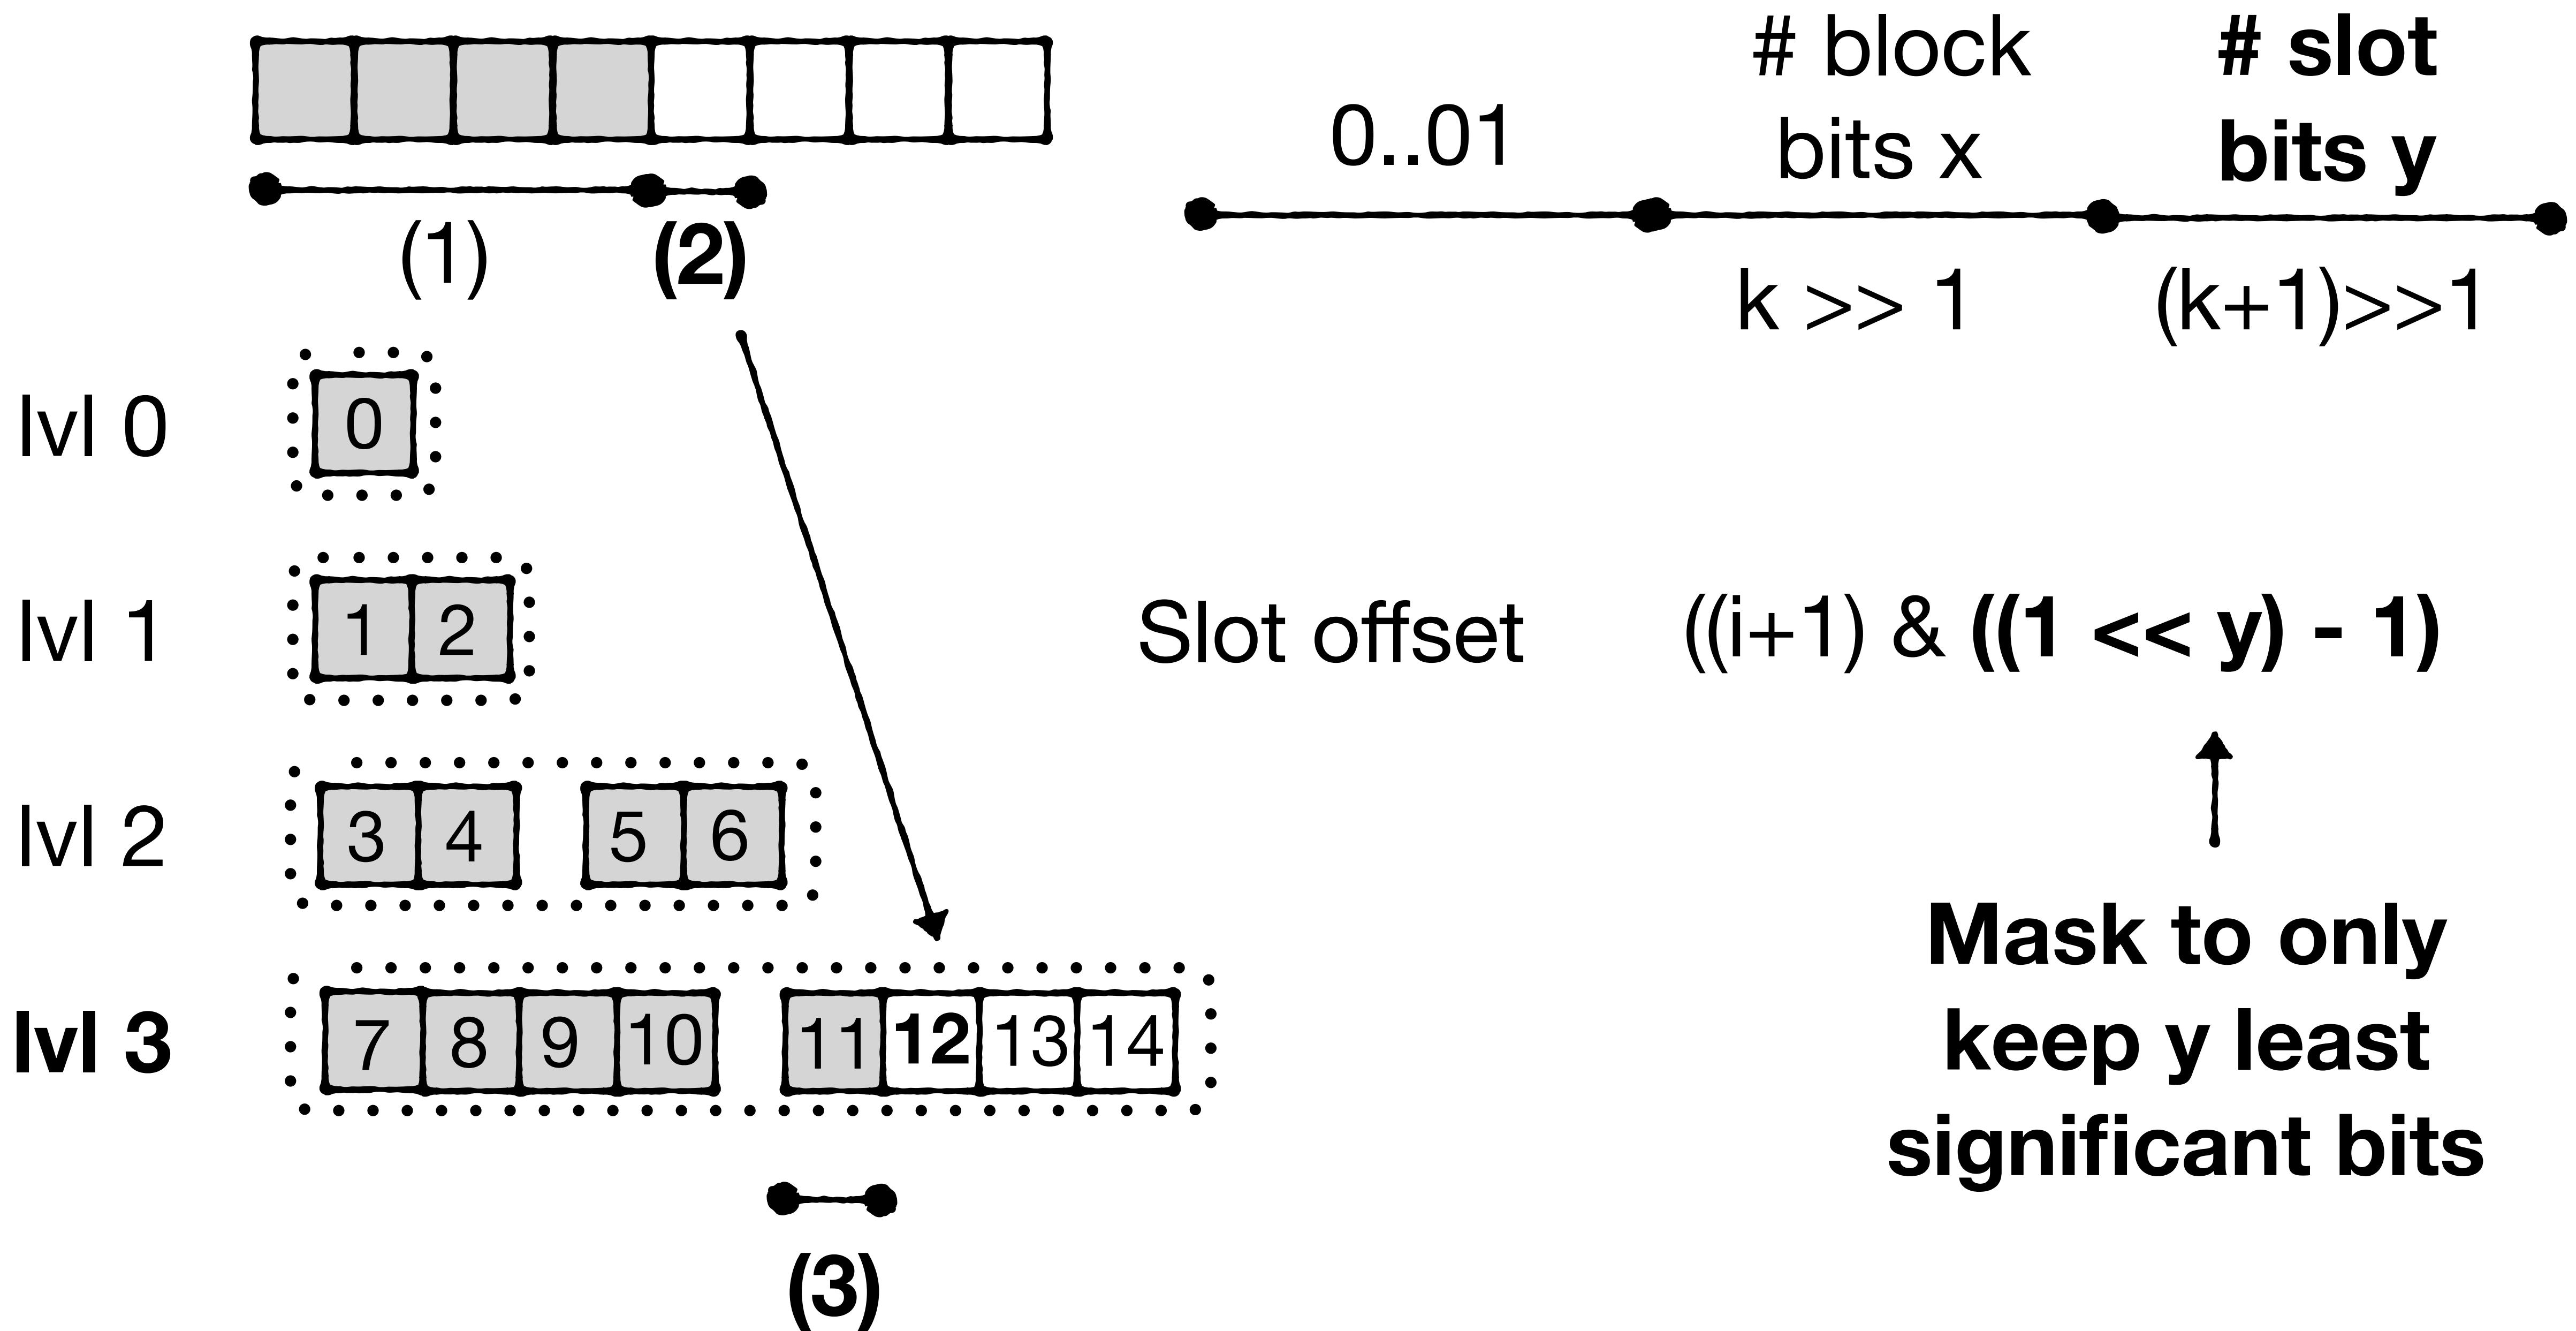


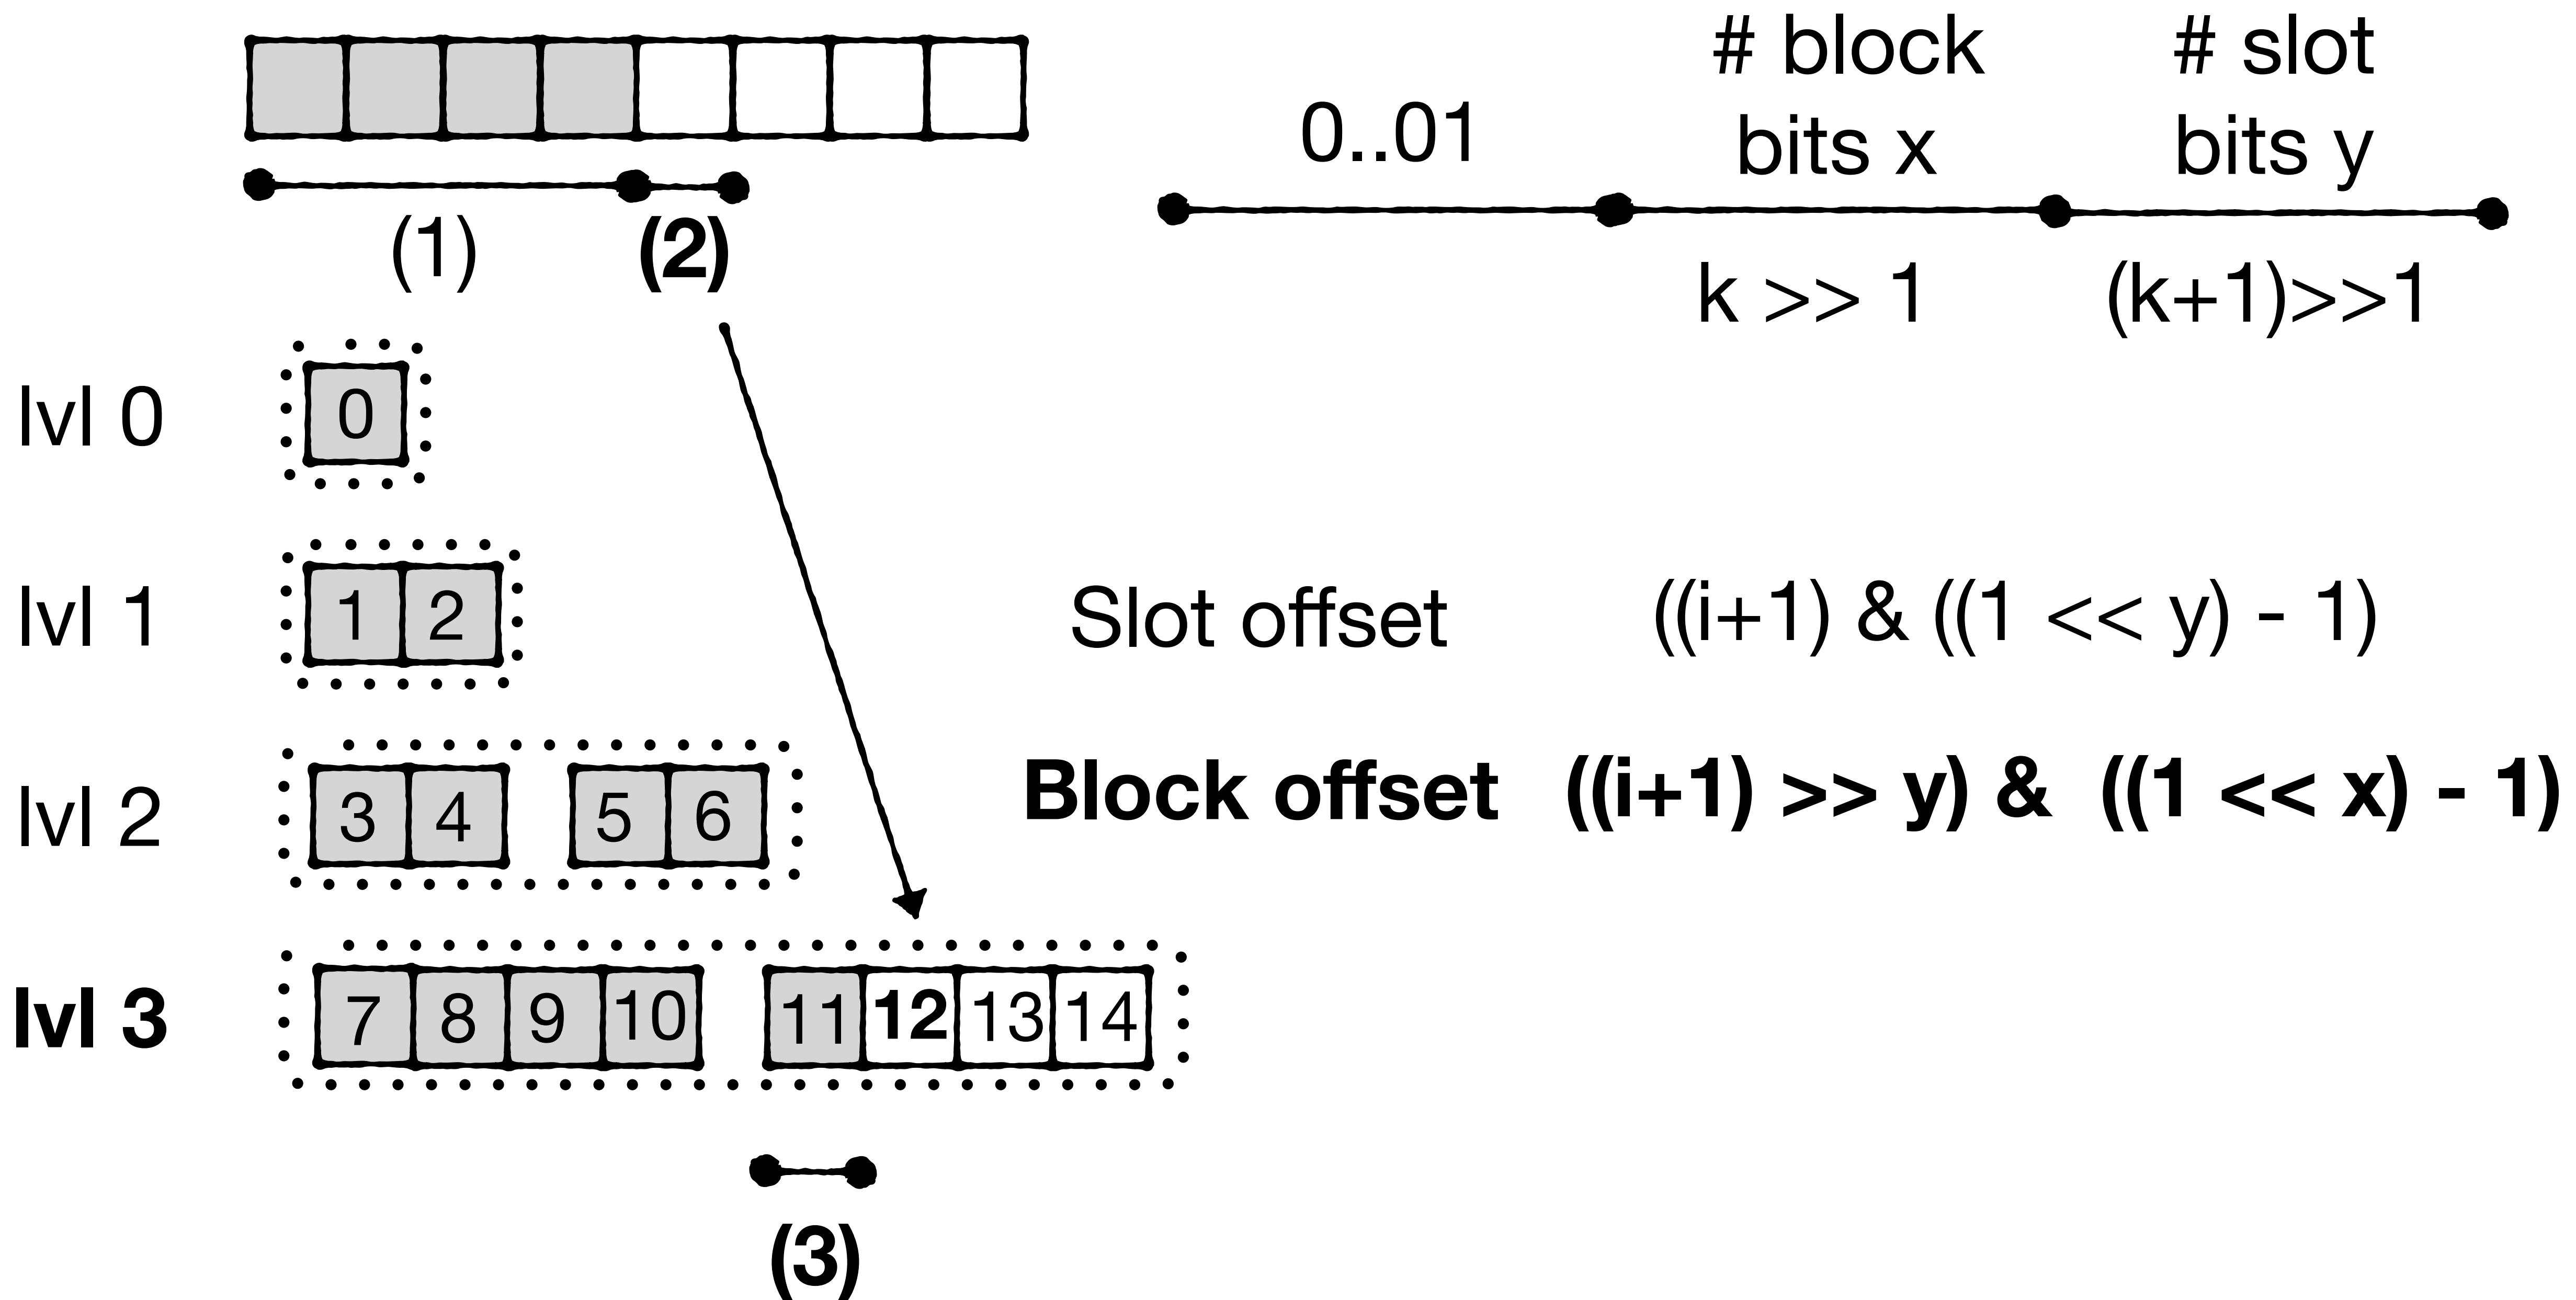


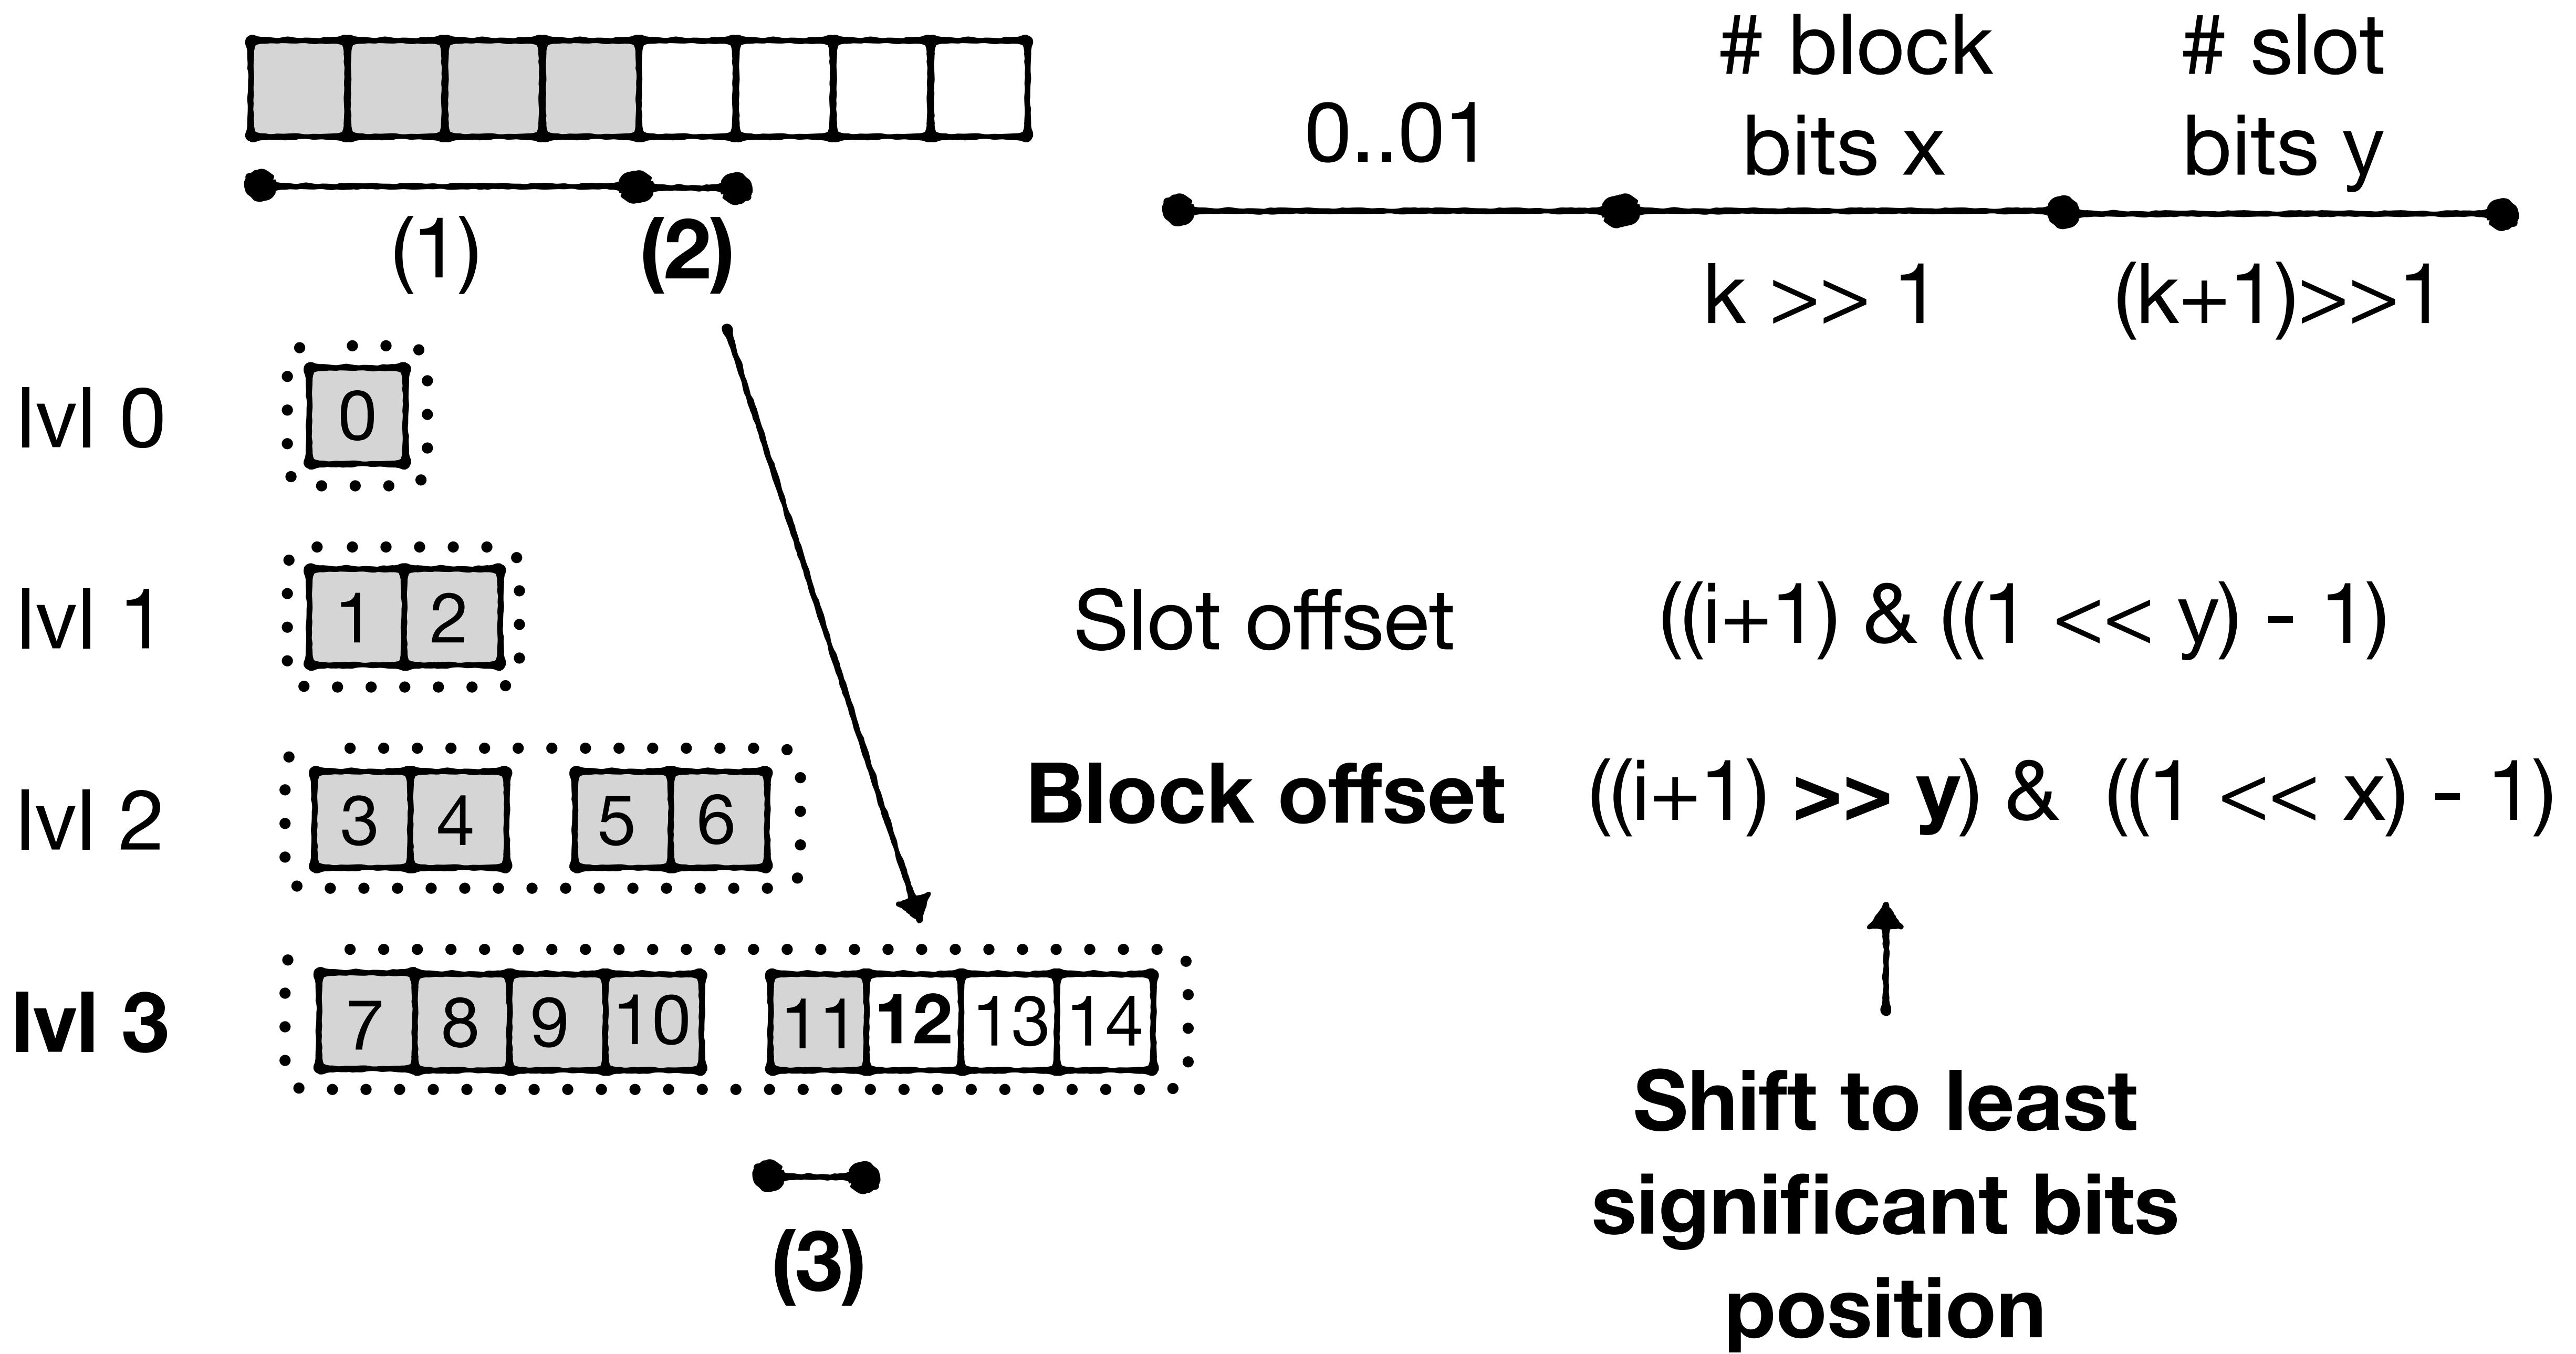


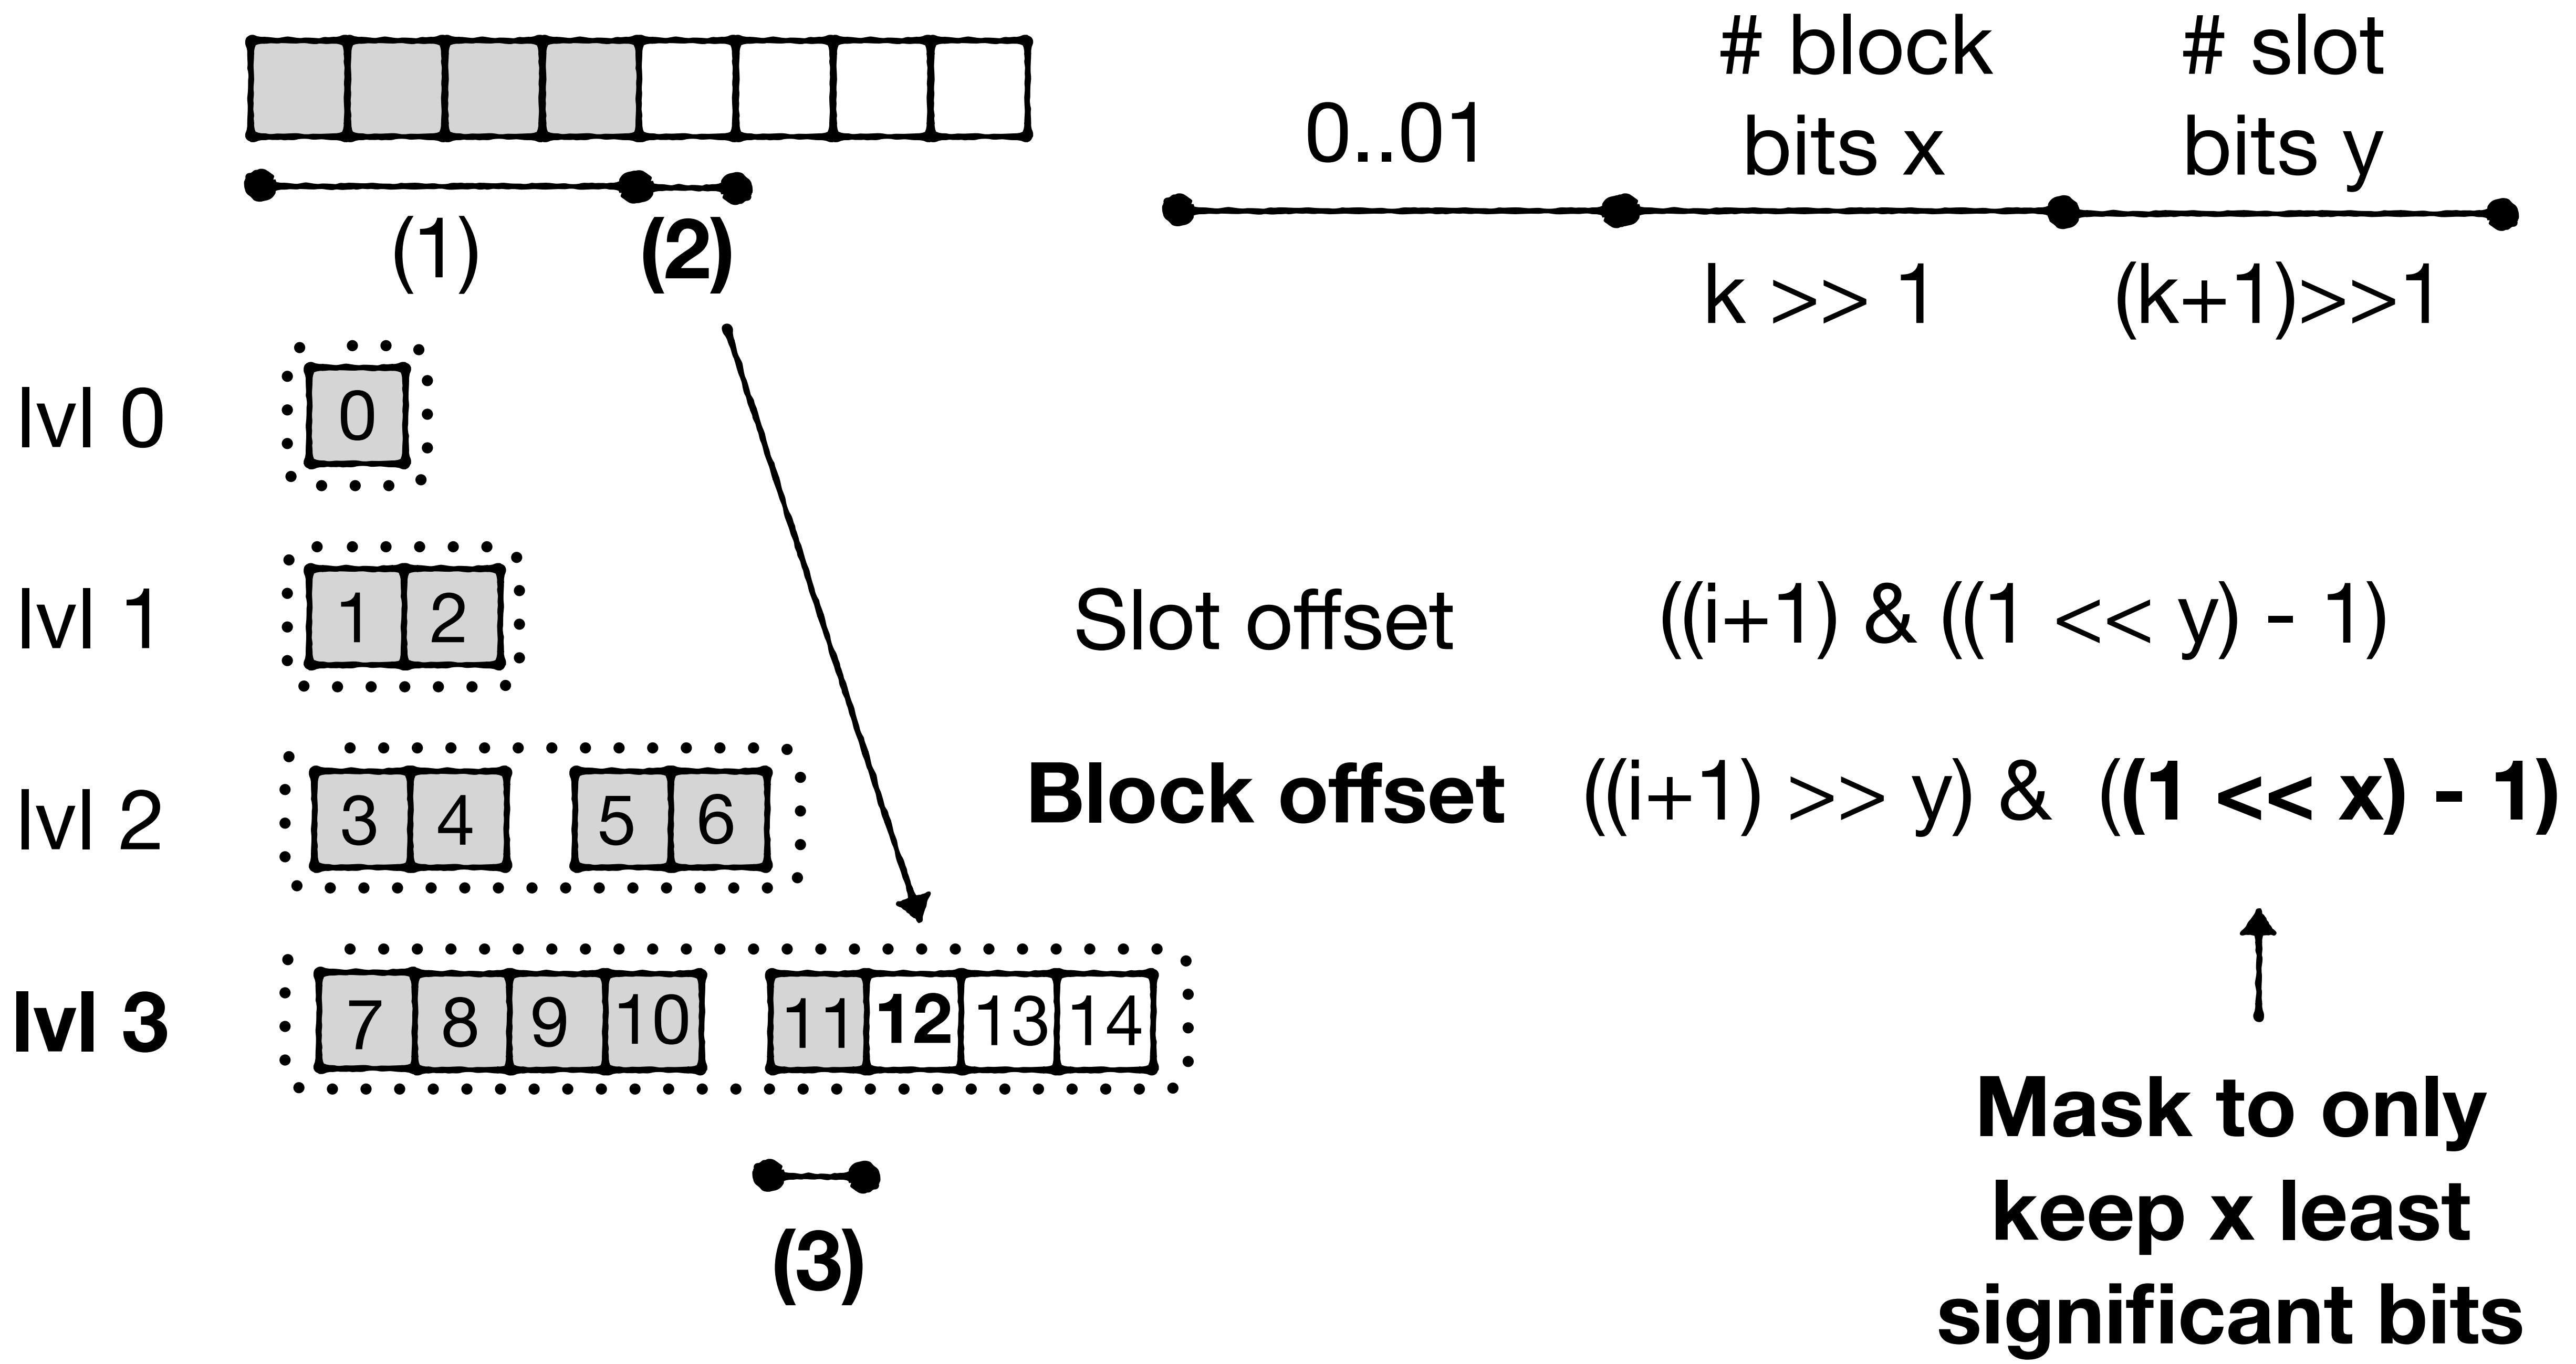


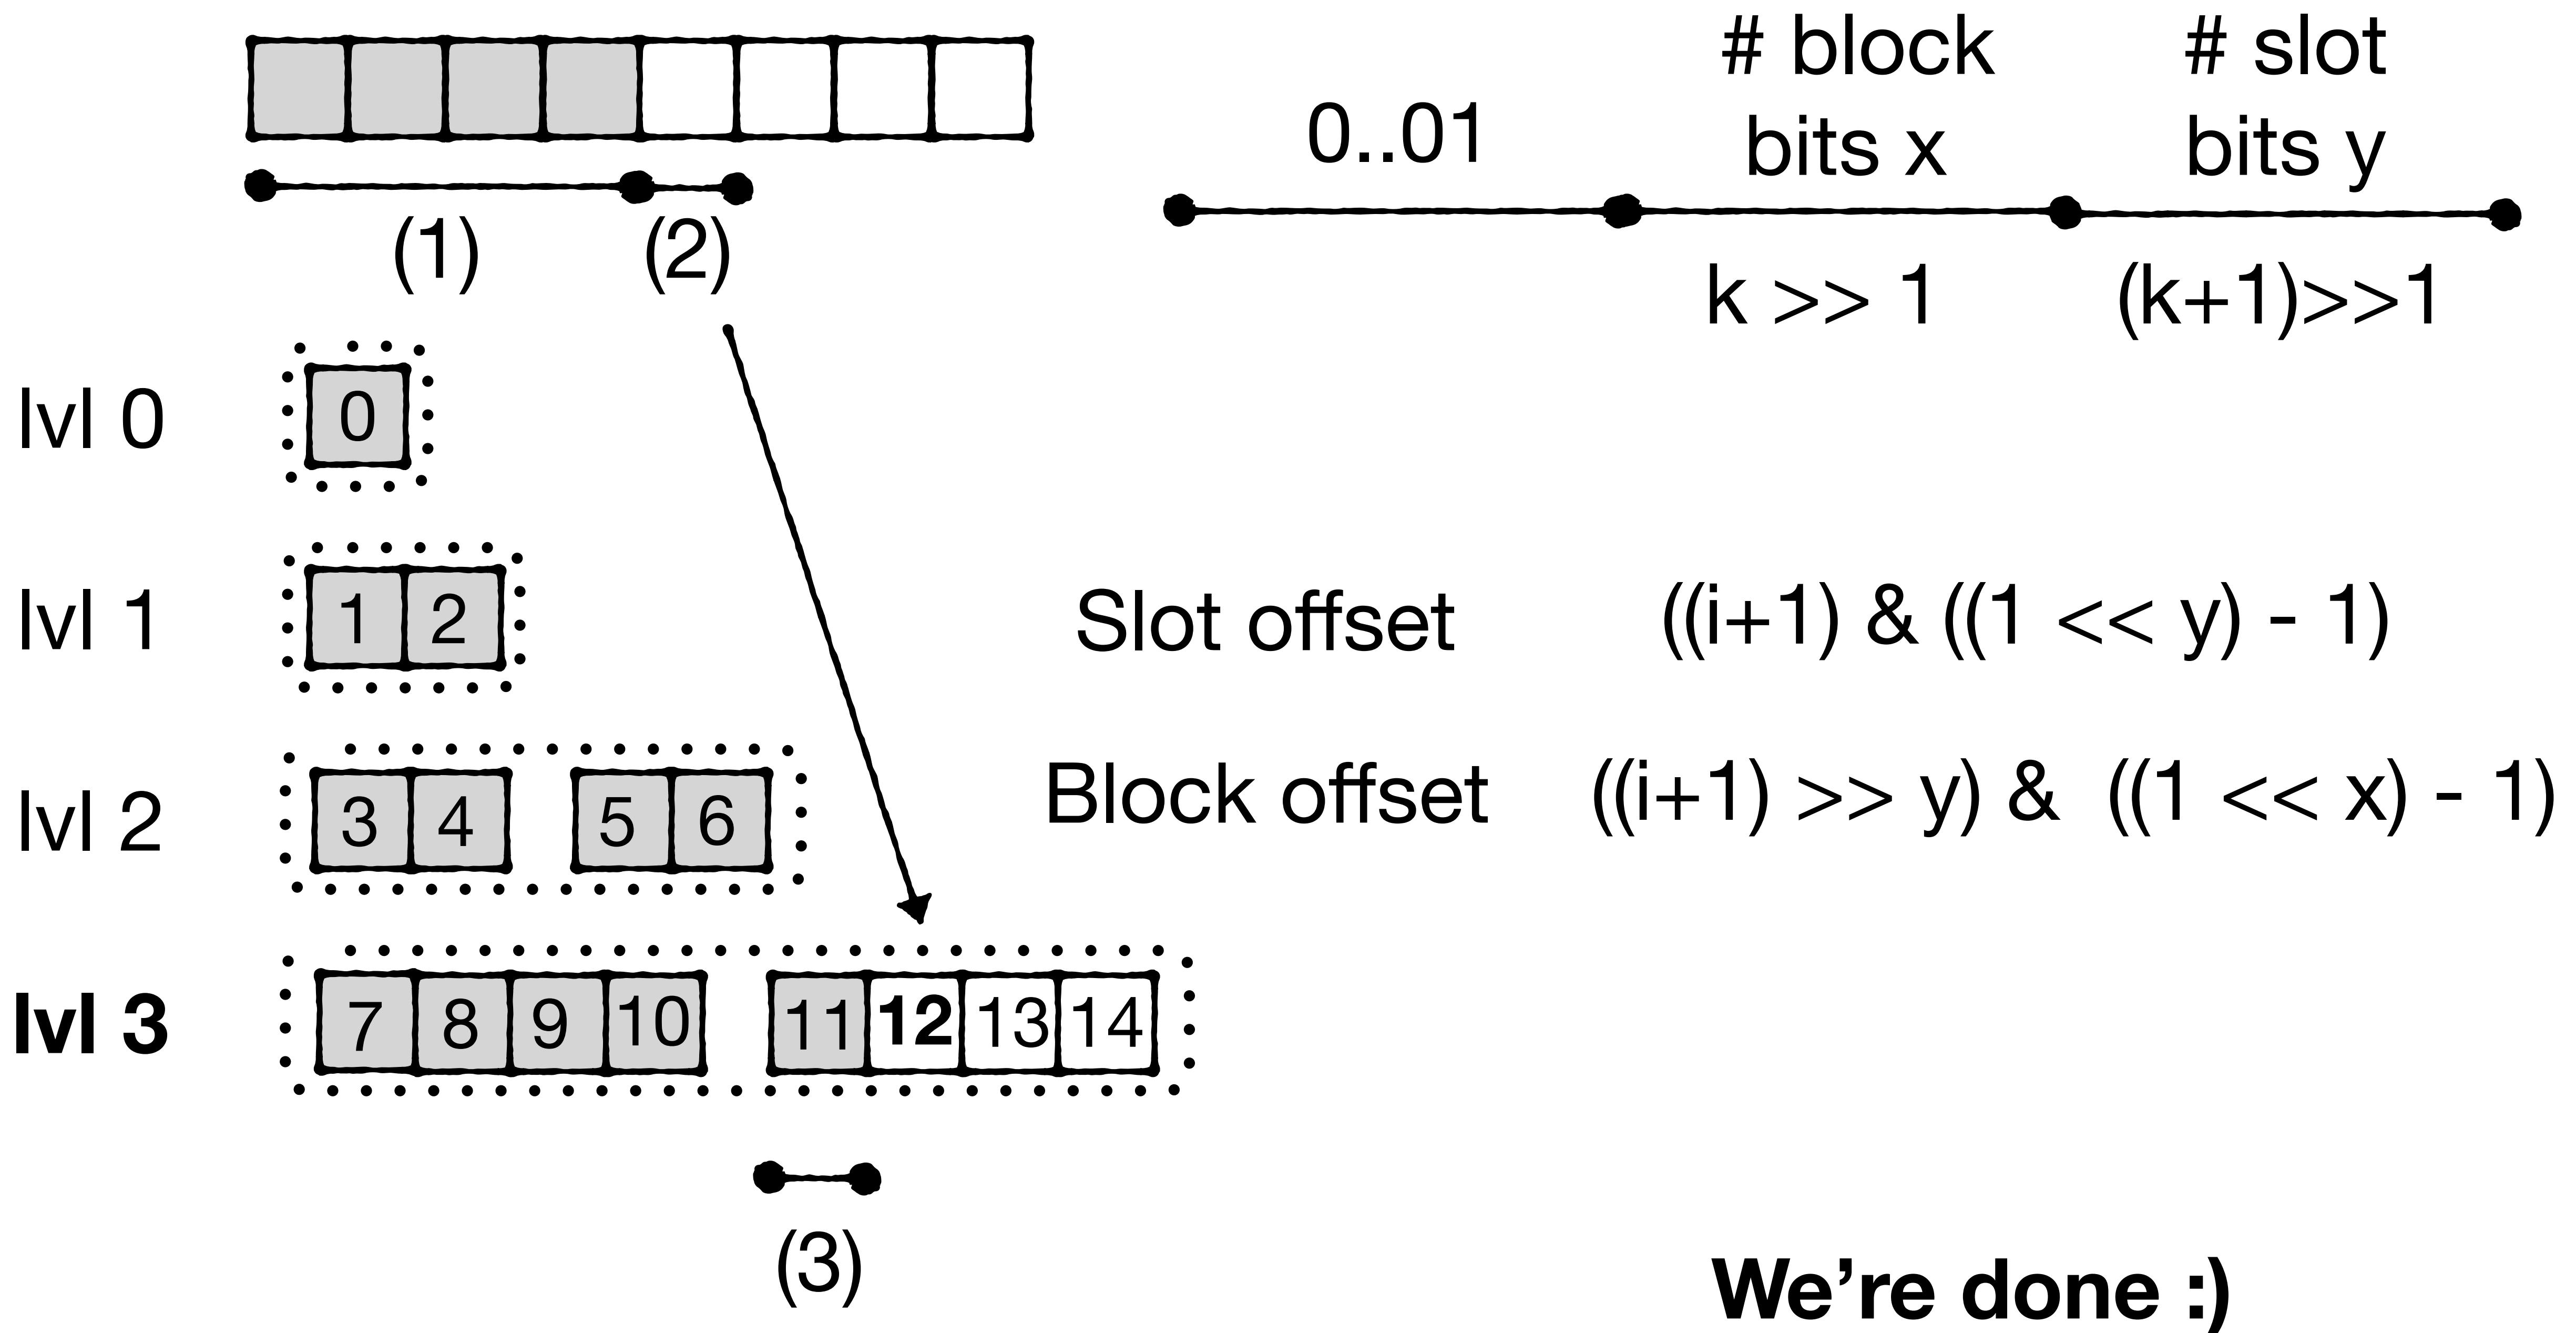












Write-amp

$O(1+N^{-0.5})$

indirection

Space-amp

$O(1+N^{-0.5})$

Read-amp

$O(1)$

Write-amp

indirection

$O(1+N^{-0.5})$

No indirection

$$\frac{G}{G - 1}$$

Space-amp

$O(1+N^{-0.5})$

$O(G)$

Read-amp

$O(1)$

1

Write-amp

indirection

$O(1+N^{-0.5})$

No indirection

$G > \phi$

$$\frac{G}{G - 1}$$

No indirection

$G < \phi$

$$\frac{G}{G - 1}$$

Space-amp

$O(1+N^{-0.5})$

$$\frac{G}{G - 1} + G$$

G to $G+1$

Read-amp

$O(1)$

1

1

Thank you :)