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?

> 10 years of research experience in data structures & algorithms for databases

In academia

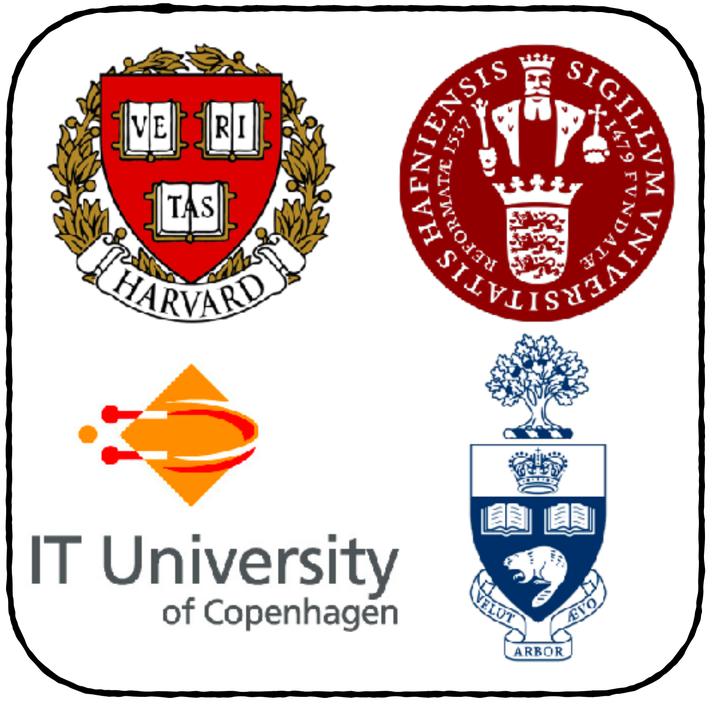

Who am I?

> 10 years of research experience in data structures & algorithms for databases

In academia

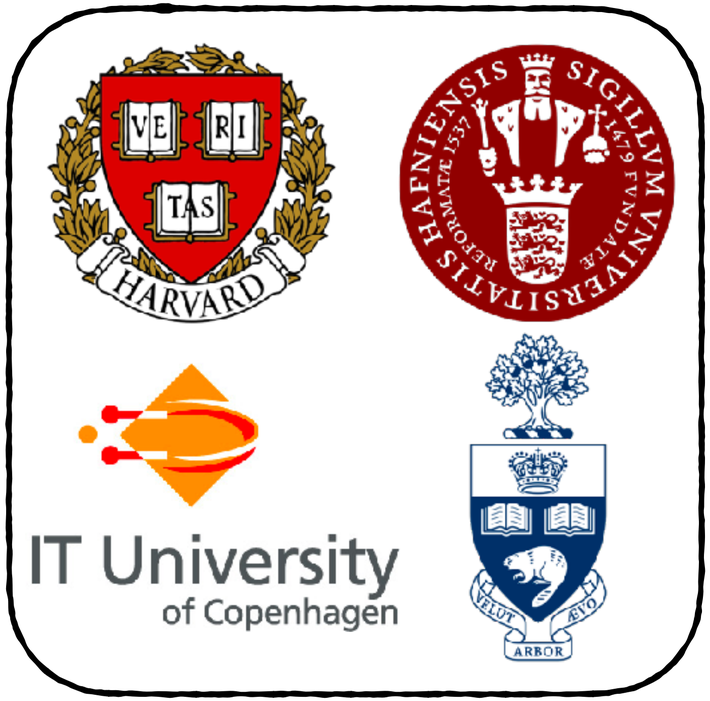

And in industry

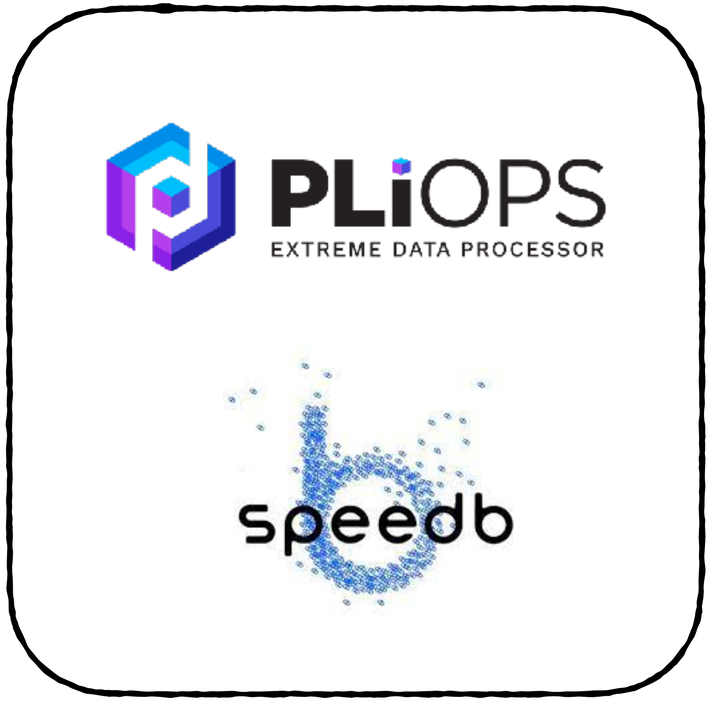

This course combines both

Theory

Practice

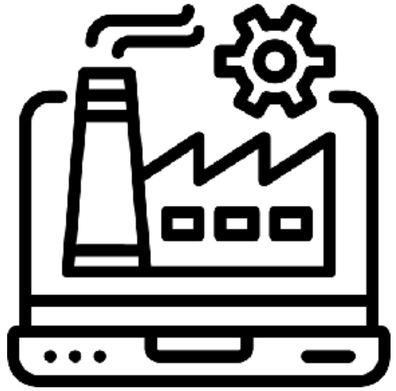

For data structures & algorithms for databases

Who are you?

Who are you?

Grad

Undergrad

Who are you? Prerequisites…

Introduction to Databases

Relational algebra, data modeling, SQL

Who are you? Prerequisites…

Introduction to Databases

Relational algebra, data modeling, SQL

Operating Systems

Concurrency & synchronization
File systems, virtual memory

Who are you? Prerequisites…

Introduction to Databases

Relational algebra, data modeling, SQL

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…

Introduction to Databases

Relational algebra, data modeling, SQL

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…

Introduction to Databases

Relational algebra, data modeling, SQL

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 important background

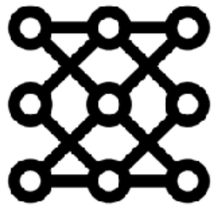

All lectures recorded

Will let you know what to catch up on

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

Seminar

Reading ≈20-30 Papers

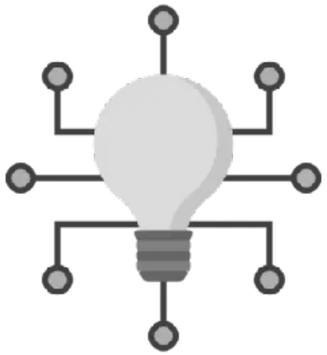

Small Research Project

Seminar

Reading ≈20-30 Papers

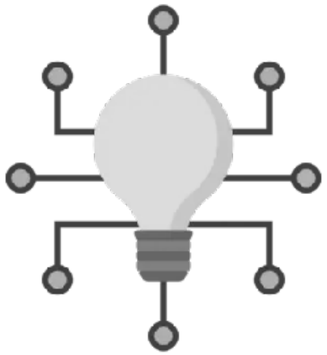

Small Research Project

Papers

topical

Classics

“Best papers”

influential (citations)

Papers

topical

Classics

“Best papers”

influential (citations)

New skill & insight

Why read papers?

Reading papers is a skill

Get research ideas

Employ the state of the art

Focus

DB is a huge field

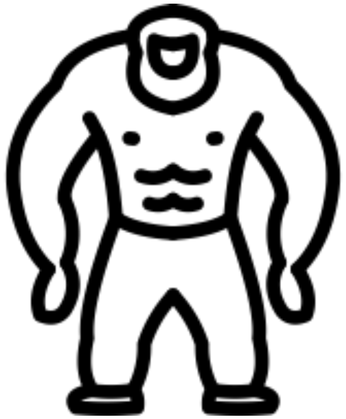

Data structures for modern hardware

For a specialized course, I may as well teach my specialty

Focus

DB is a huge field

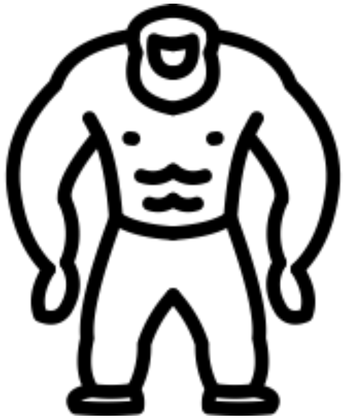

Data structures for modern hardware

For a specialized course, I may as well teach my specialty

Focus

DB is a huge field

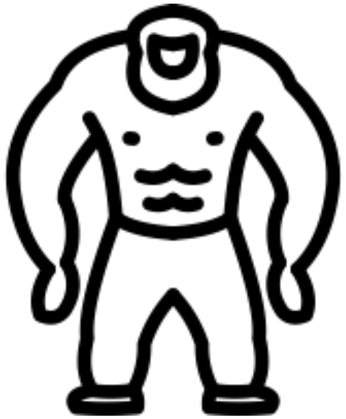

Data structures for modern hardware

For a specialized course, I may as well teach my specialty

Website

CSC2525

Research Topics in Database Management

Instructor: Niv Dayan
Lectures: Wednesday 13:00-15:00 (BA1200)
Office Hours: after each class

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

12 Class Sessions

12 Class Sessions

First 5 ish are lectures by me

12 Class Sessions

First 5 ish are lectures by me

Subsequent 7 ish

1-1.5 hour lecture

Student Presentation

12 Class Sessions

First 5 ish are lectures by me

Subsequent 7 ish

1-1.5 hour lecture

Student Presentation

More on presentations later

Use the first two lectures wisely

Dynamic Arrays & Filter Data Structures

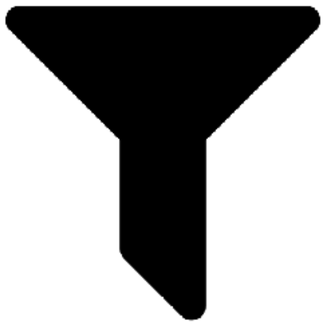

Enjoy the material? you’re in the right place

Use the first two lectures wisely

Dynamic Arrays &
Filter Data
Structures

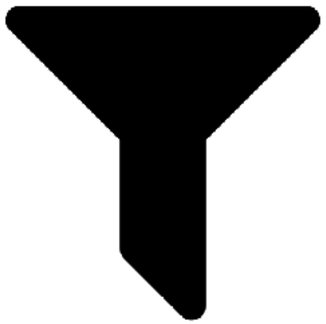

**Enjoy the material?
you're in the right
place**

Participation

You are required to attend each class

Read papers in advance

Participate in class discussions

Project

Implement & evaluate

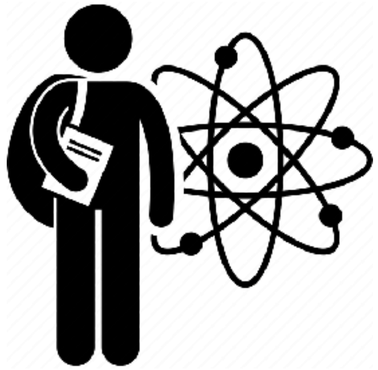

Proposals due by mid-Feb

Project

Implement & evaluate

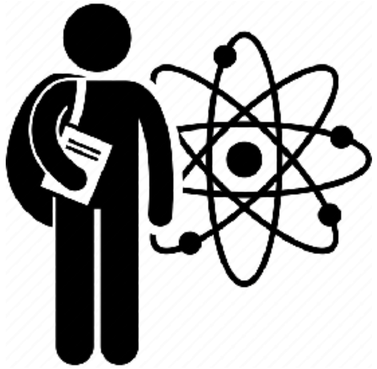

Proposals due by mid-Feb

More on this later

Oral Exam

20-30 min per student

Papers & Material

Grade Components

(1) Project Report & Code

(2) Paper presentation

(3) Oral exam

Grade Components

(1) Project Report & Code

(2) Paper presentation

(3) Oral exam

Precise breakdown to be announced later

Office Hours

Right after class

You’re encouraged to reach out before your presentation

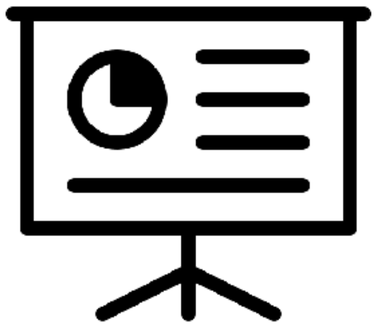

Post questions for everyone’s benefit!

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

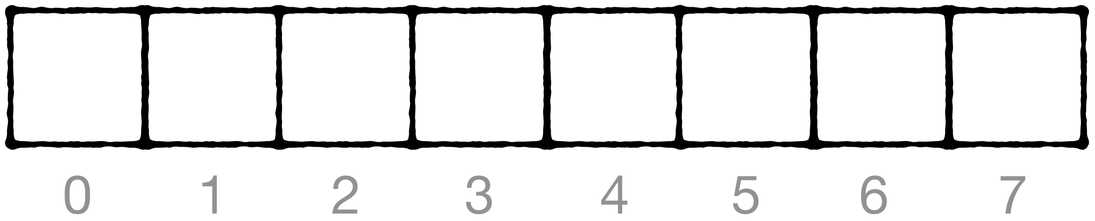

Arrays

Fixed width slots

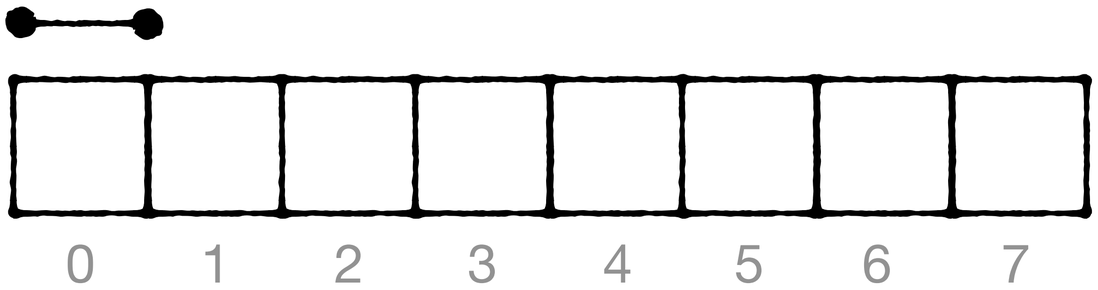

Fixed width slots

e.g., integers or floating points

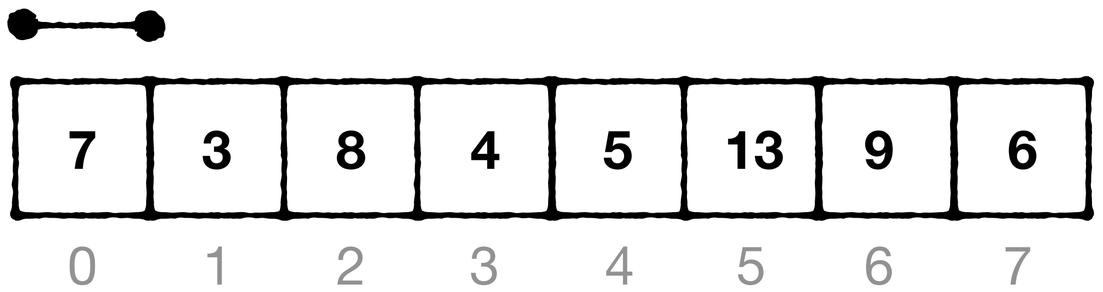

Fixed width slots

Or pointers to complex types

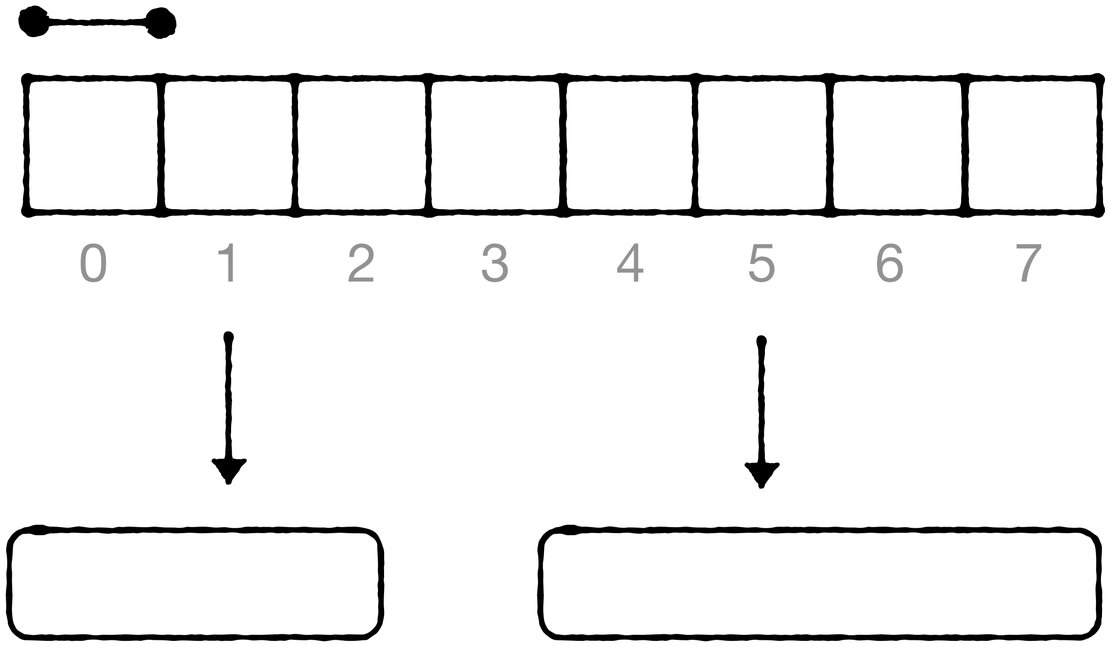

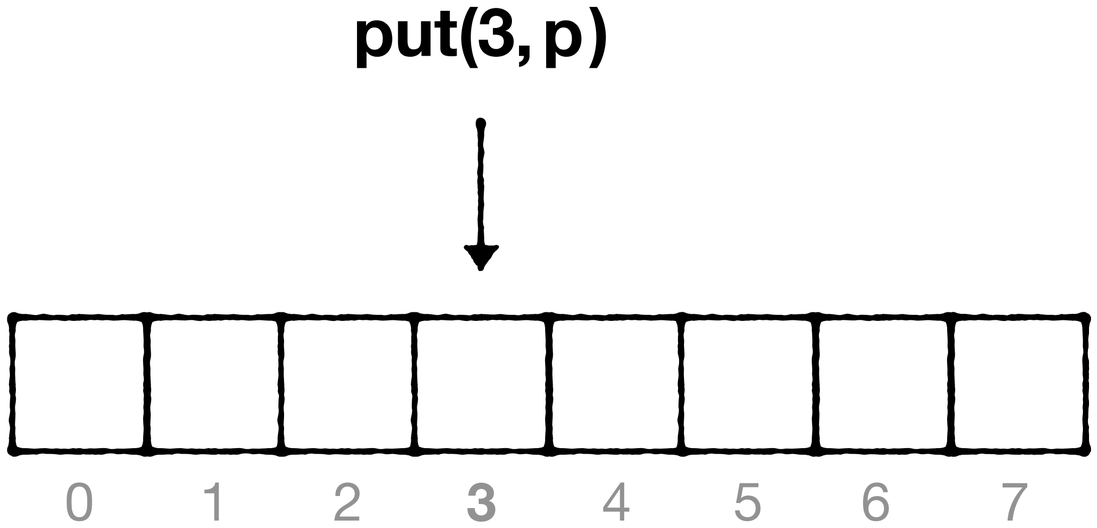

Supports random access

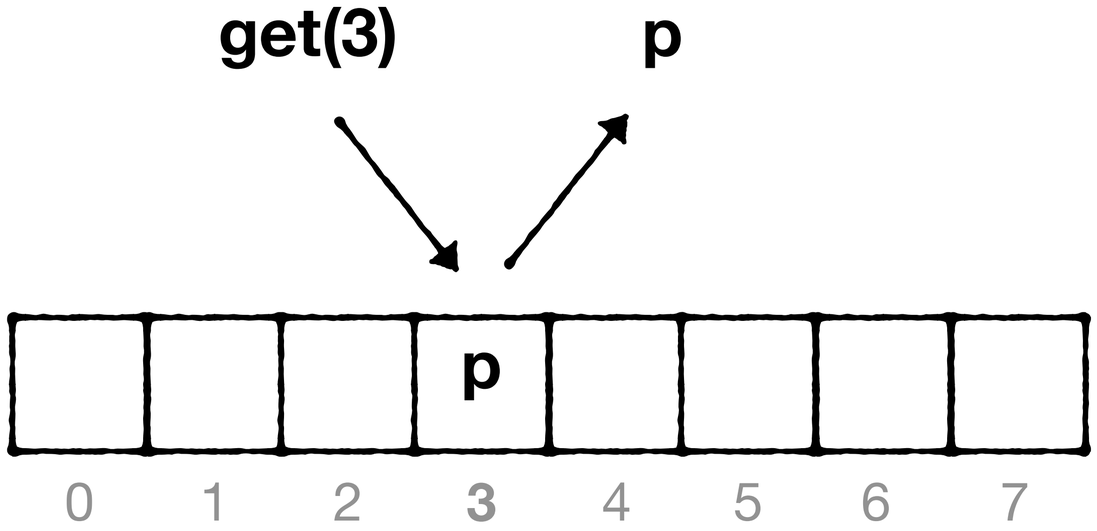

Supports random access

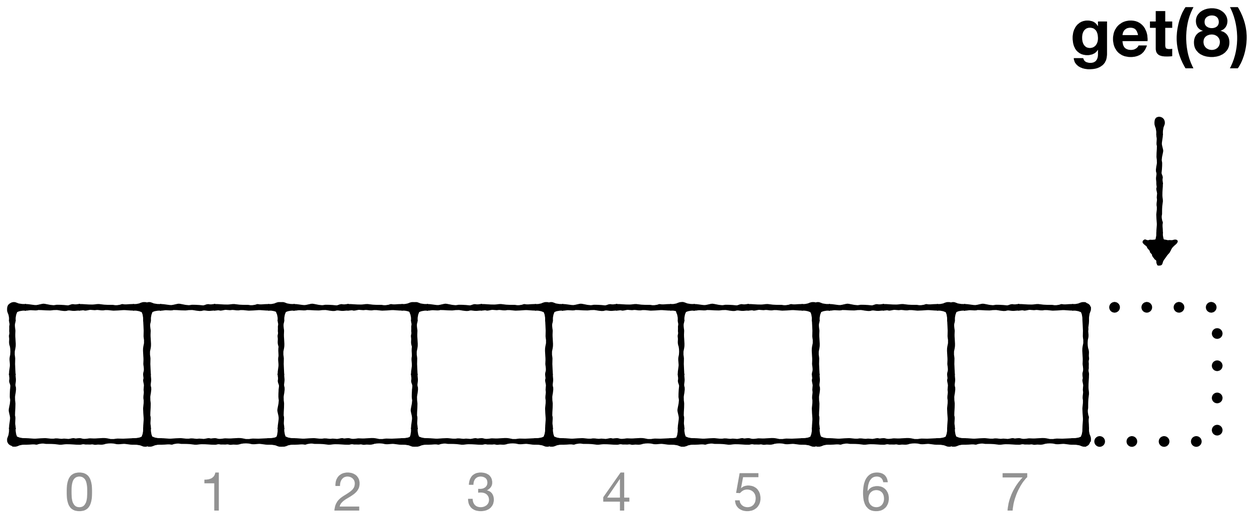

Overflow error (e.g., java)

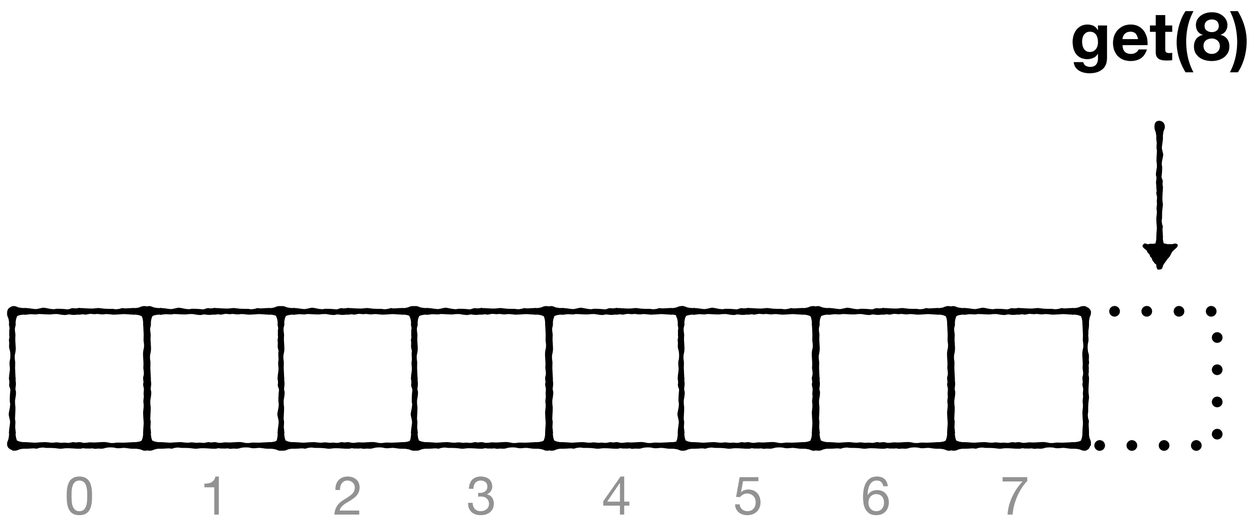

Overflow error (e.g., java)

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

How to keep inserting when out of space?

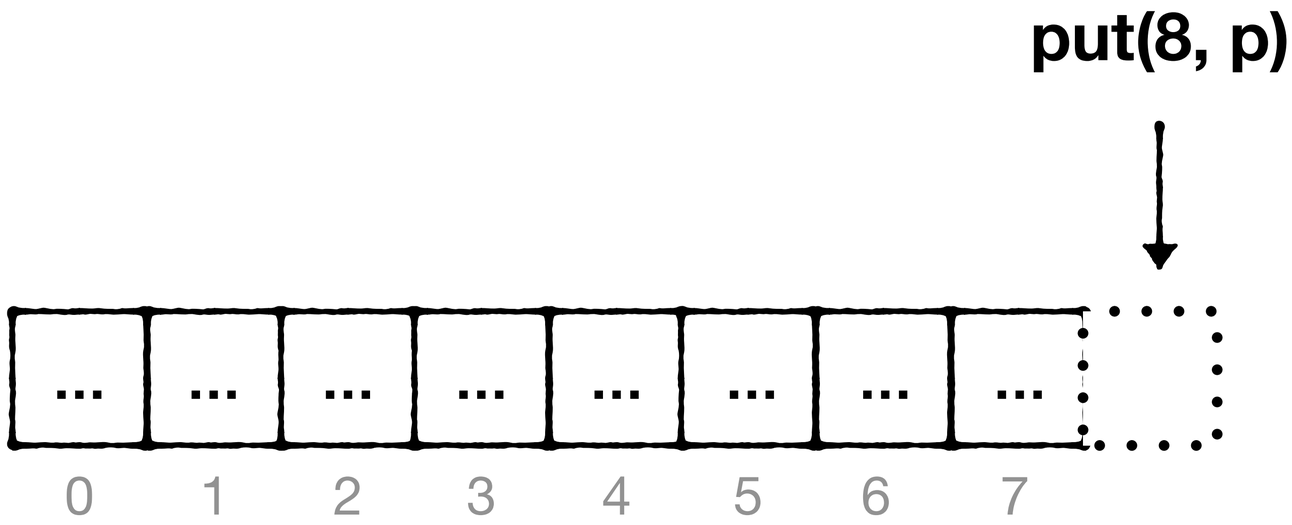

How to keep inserting when out of space?

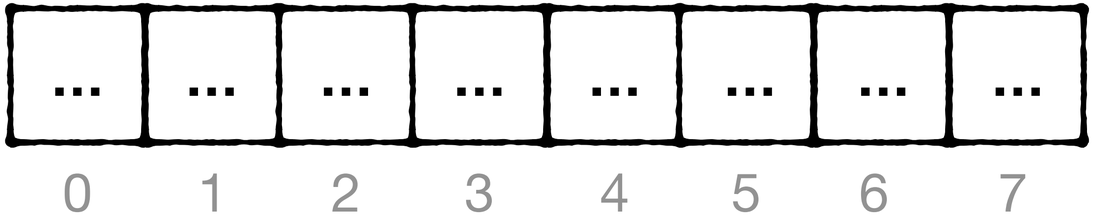

Allocate

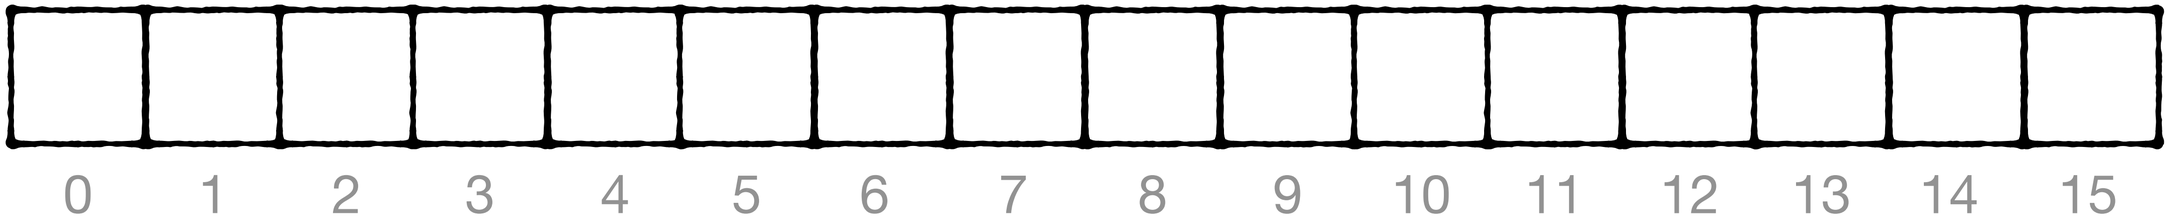

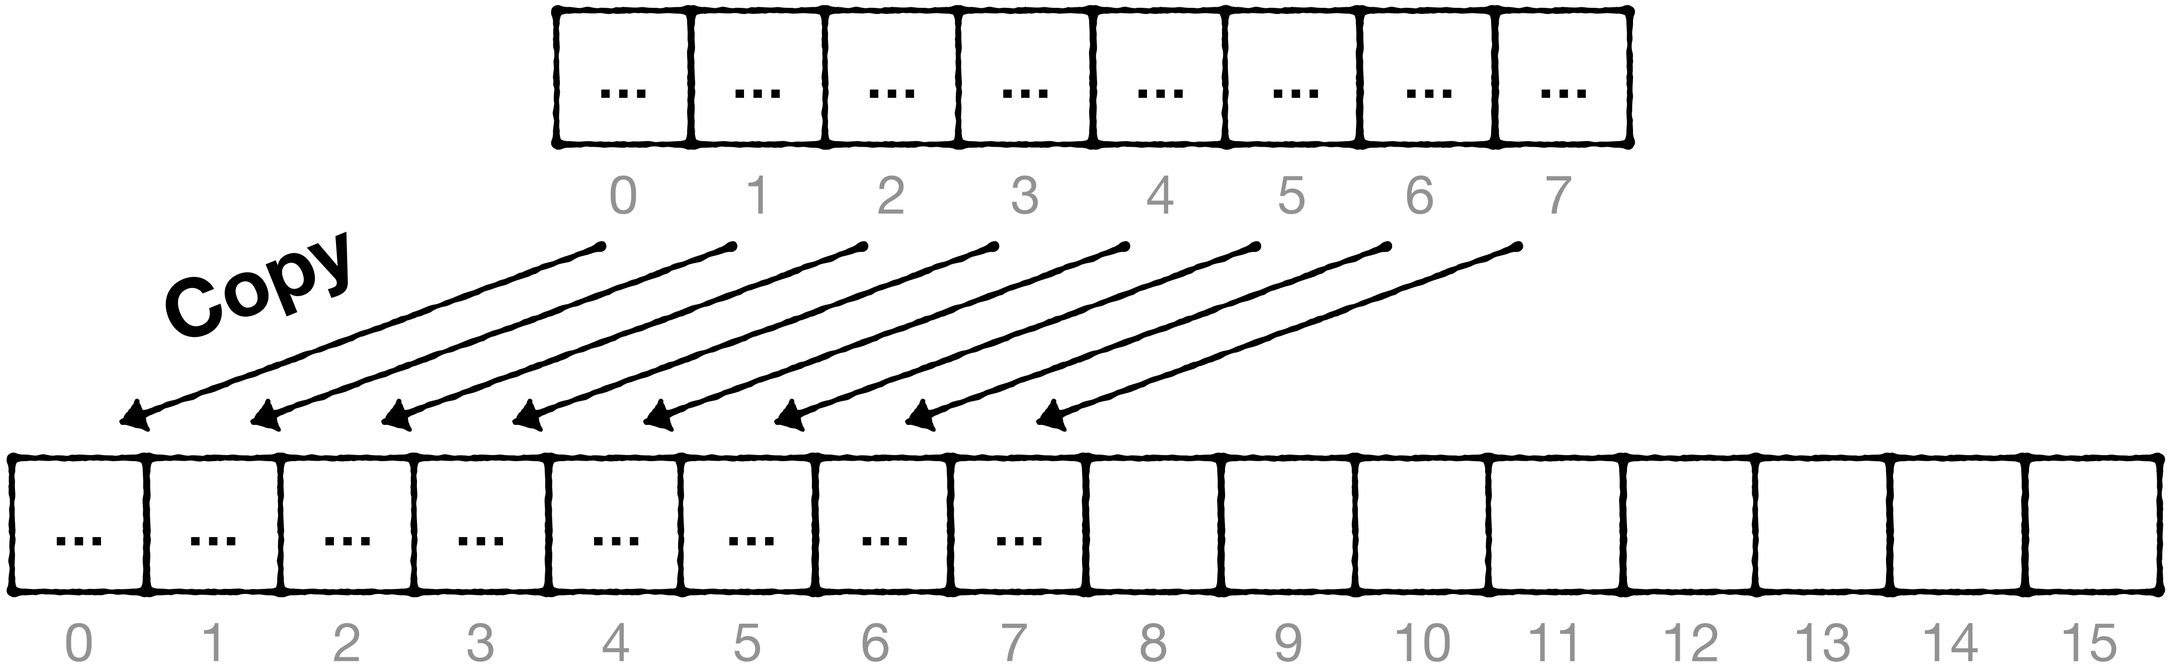
...
...
...
...
...
...
...
...
0
1
2
3
4
5
6
7
Copy
...
...
...
...
...
...
...
...
0
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15

Deallocate

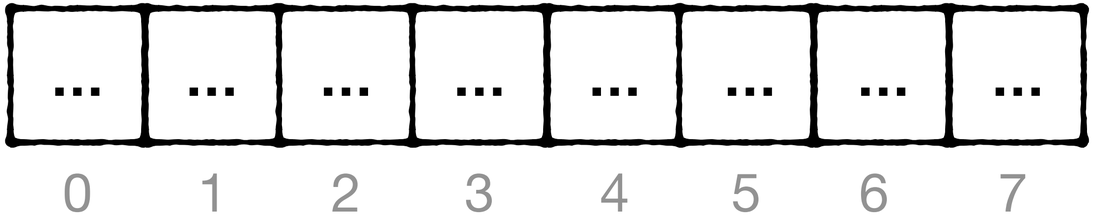

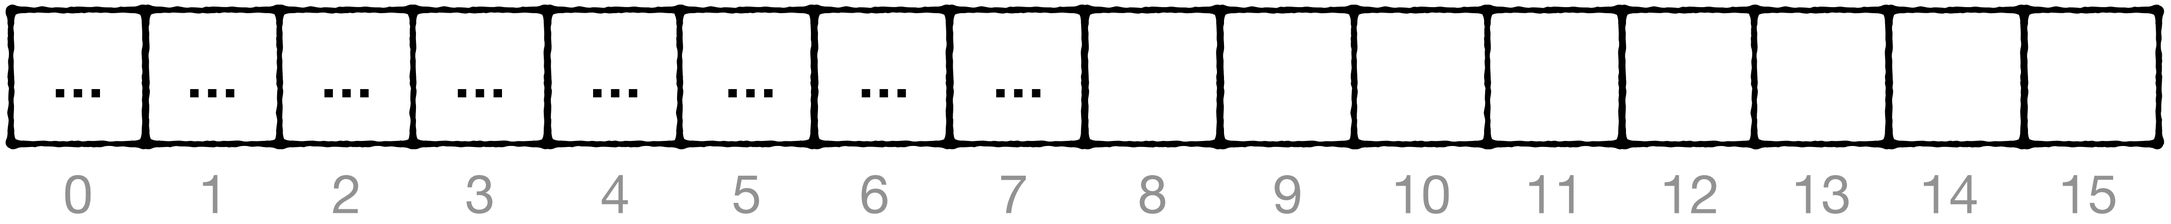

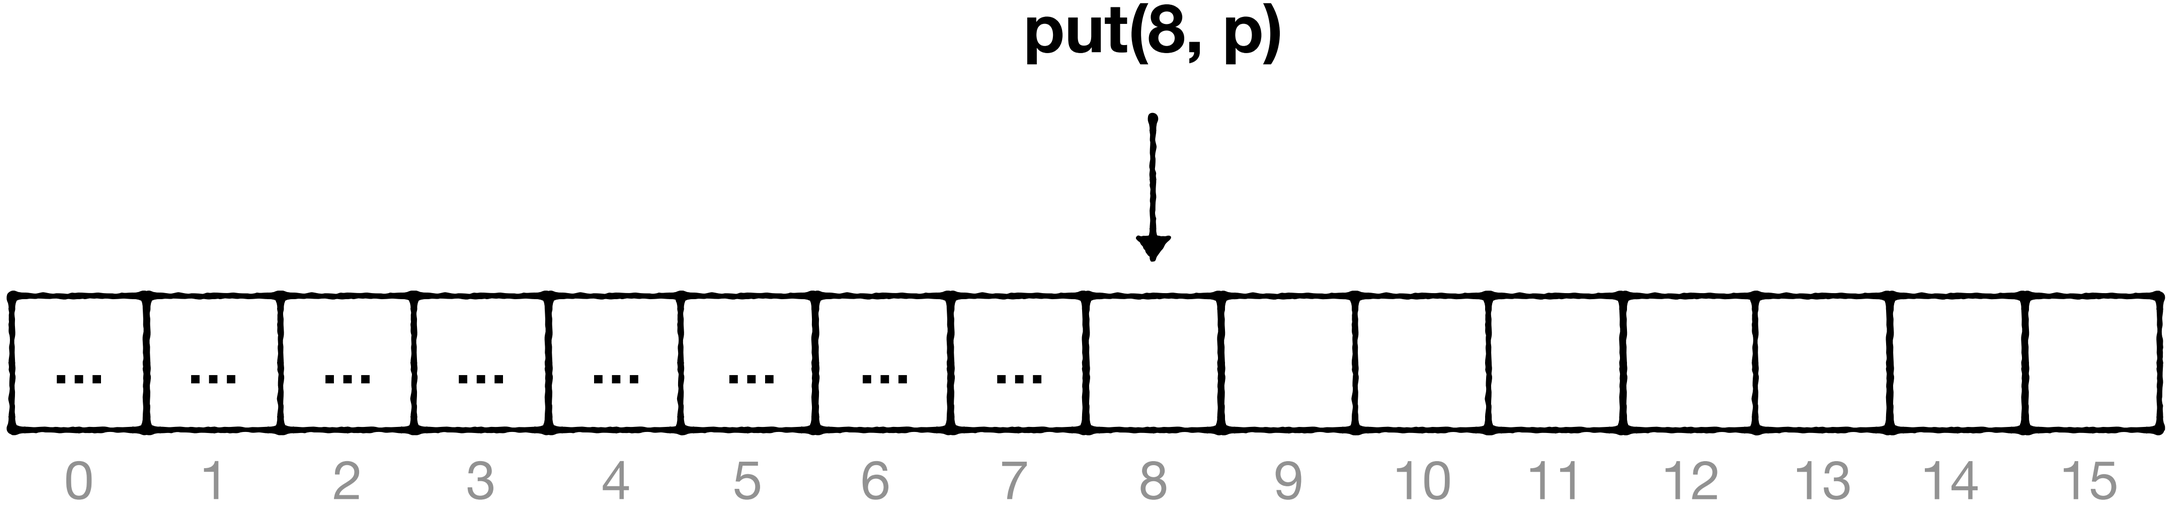
put(8, p)
...
...
...
...
...
...
...
...
0
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15

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;
}
```

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;
}
```

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:end;
}
```

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

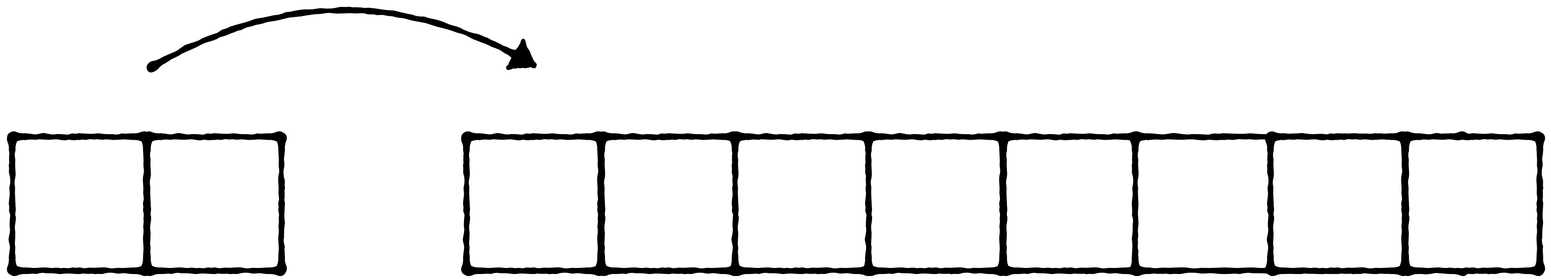

What if Growth Factor G is **too high**

e.g., 4

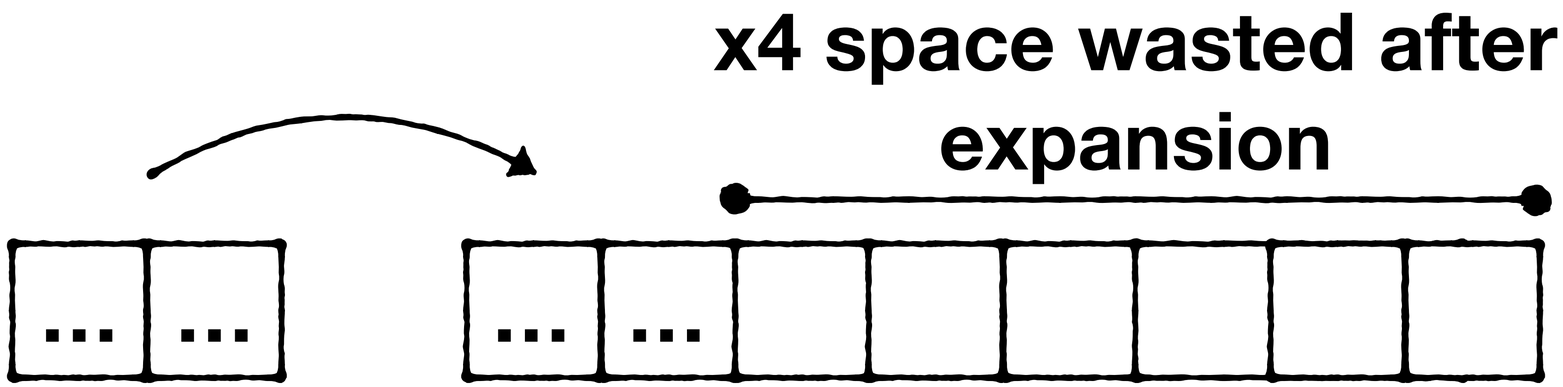

What if Growth Factor G is **too high**

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

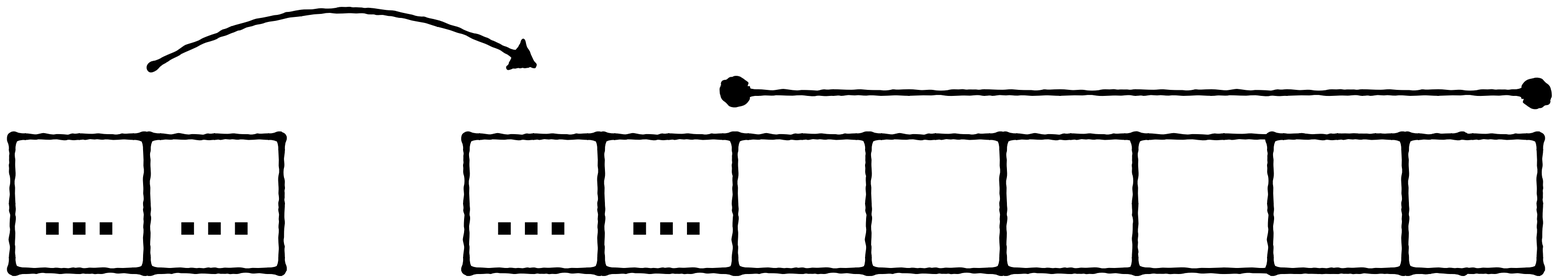

What if Growth Factor G is **too high**

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

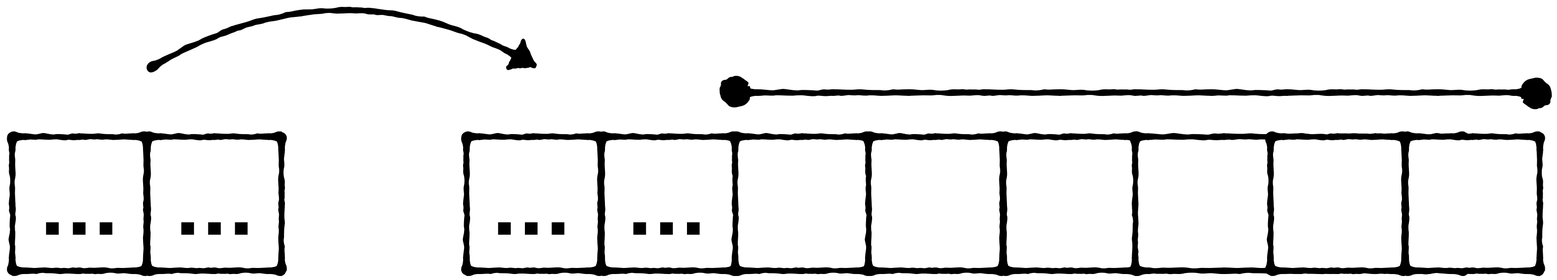

What if Growth Factor G is **too high**

Space-amplification = G

Right after expansion

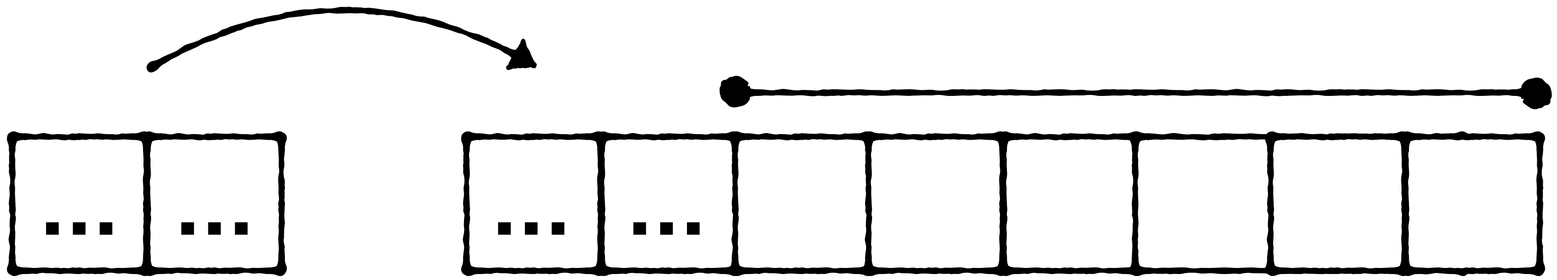

What if Growth Factor G is **too high**

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

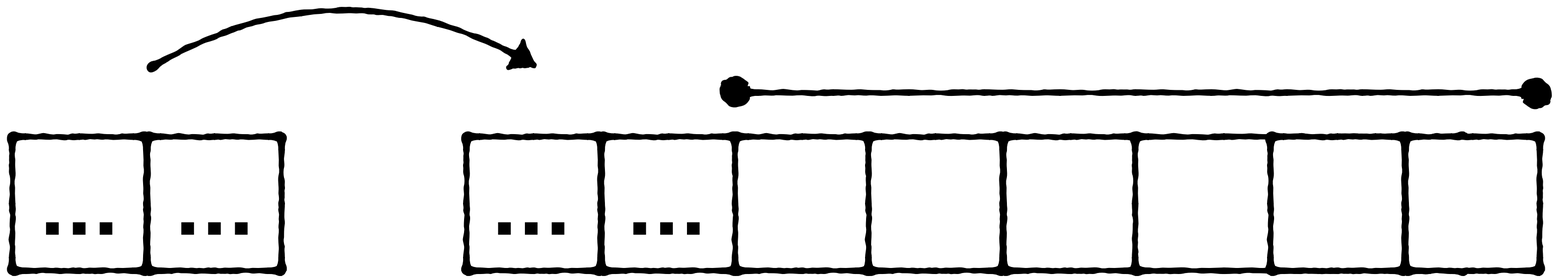

What if Growth Factor G is **too low**

e.g., 1.2

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

e.g., 1.2

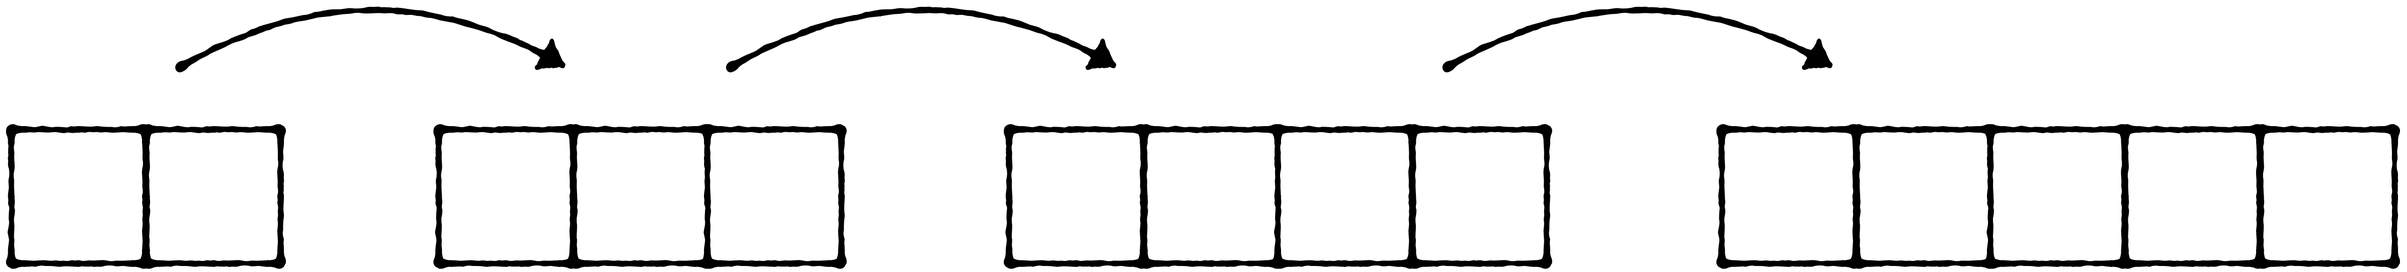

What if Growth Factor G is **too low**

e.g., 1.2

Insertion overheads increase

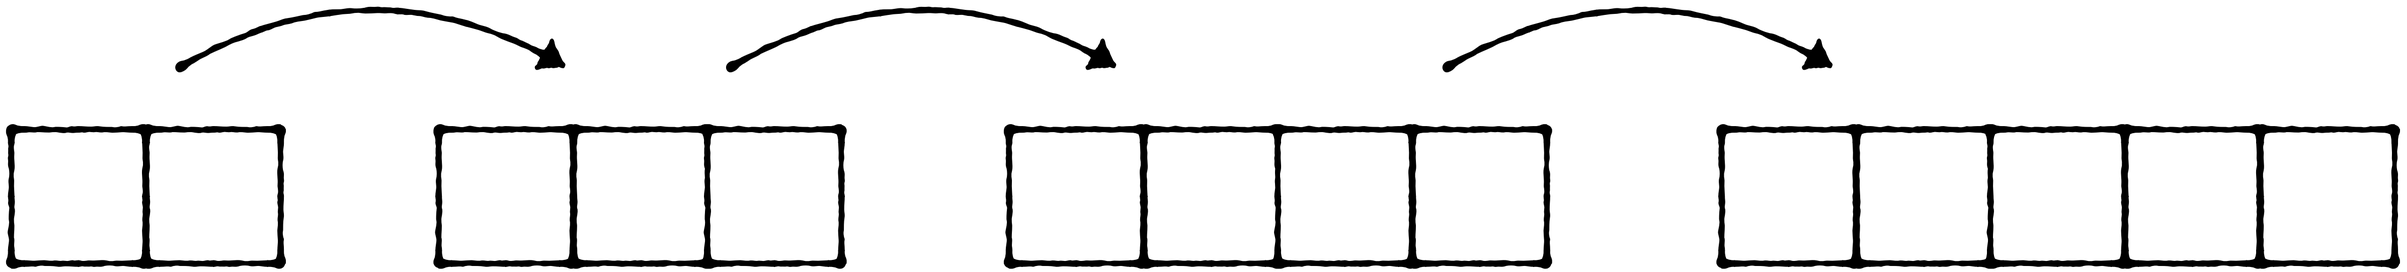

What if Growth Factor G is **too low**

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

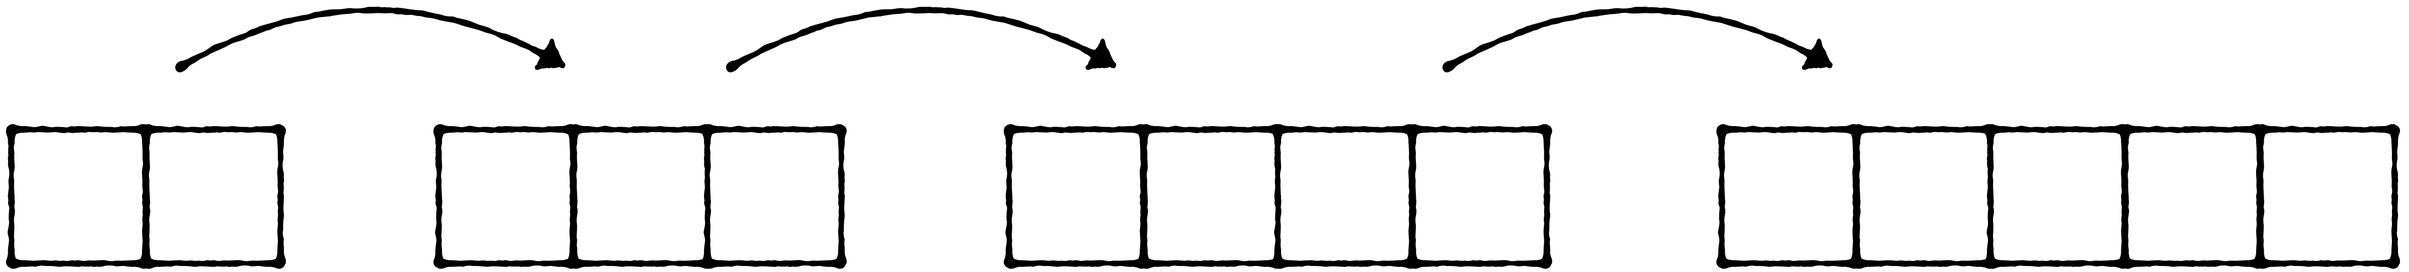

What if Growth Factor G is **too low**

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

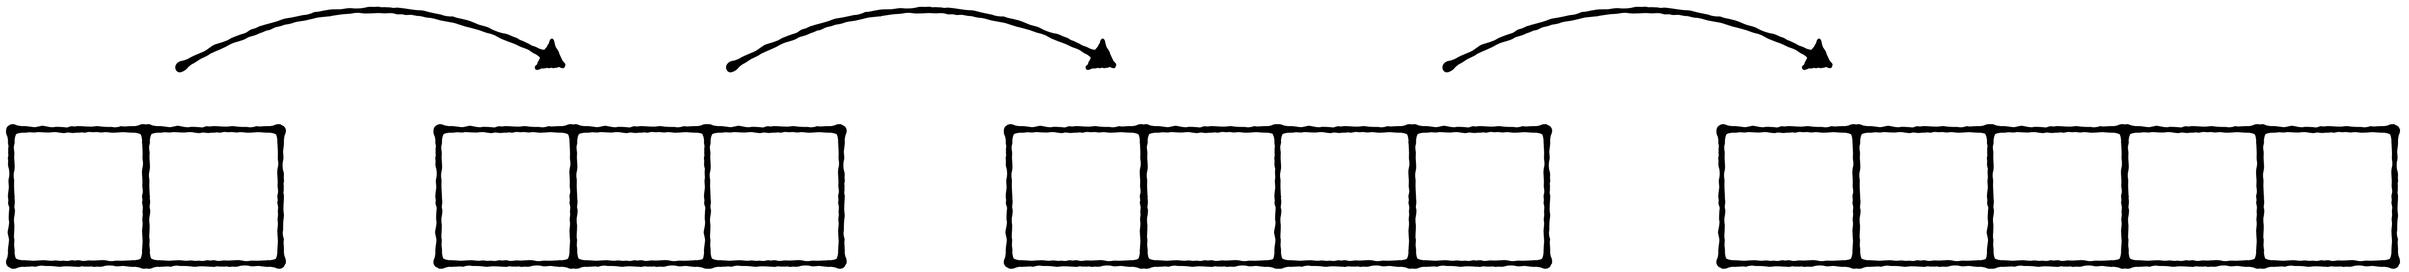

What if Growth Factor G is **too low**

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

Geometric series sum

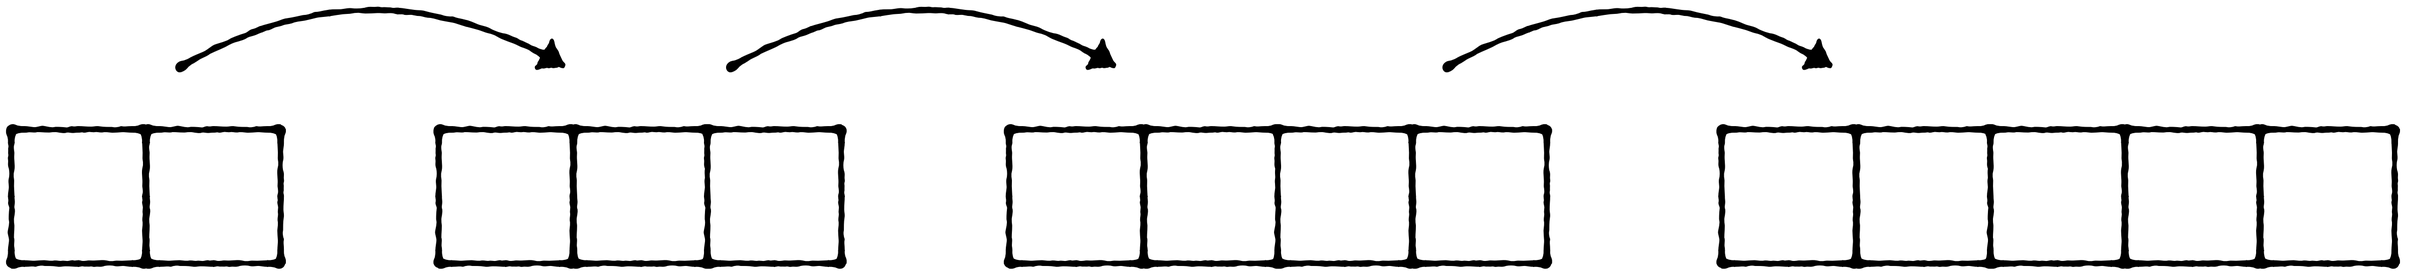

Growth factor G impact

Space-amplification

$$G+1$$

Write-amplification

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

Growth factor G impact

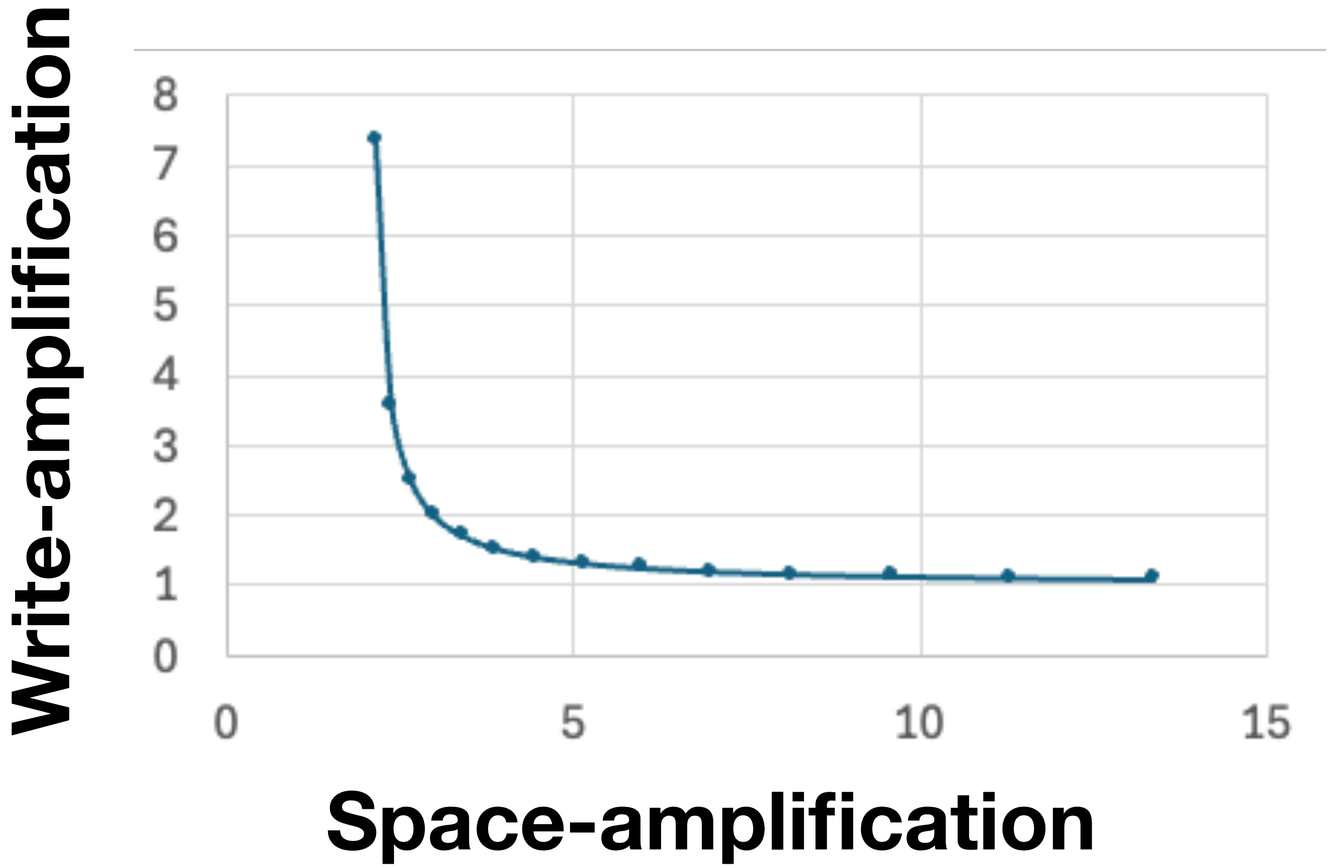

Growth factor 2 achieves a good balance

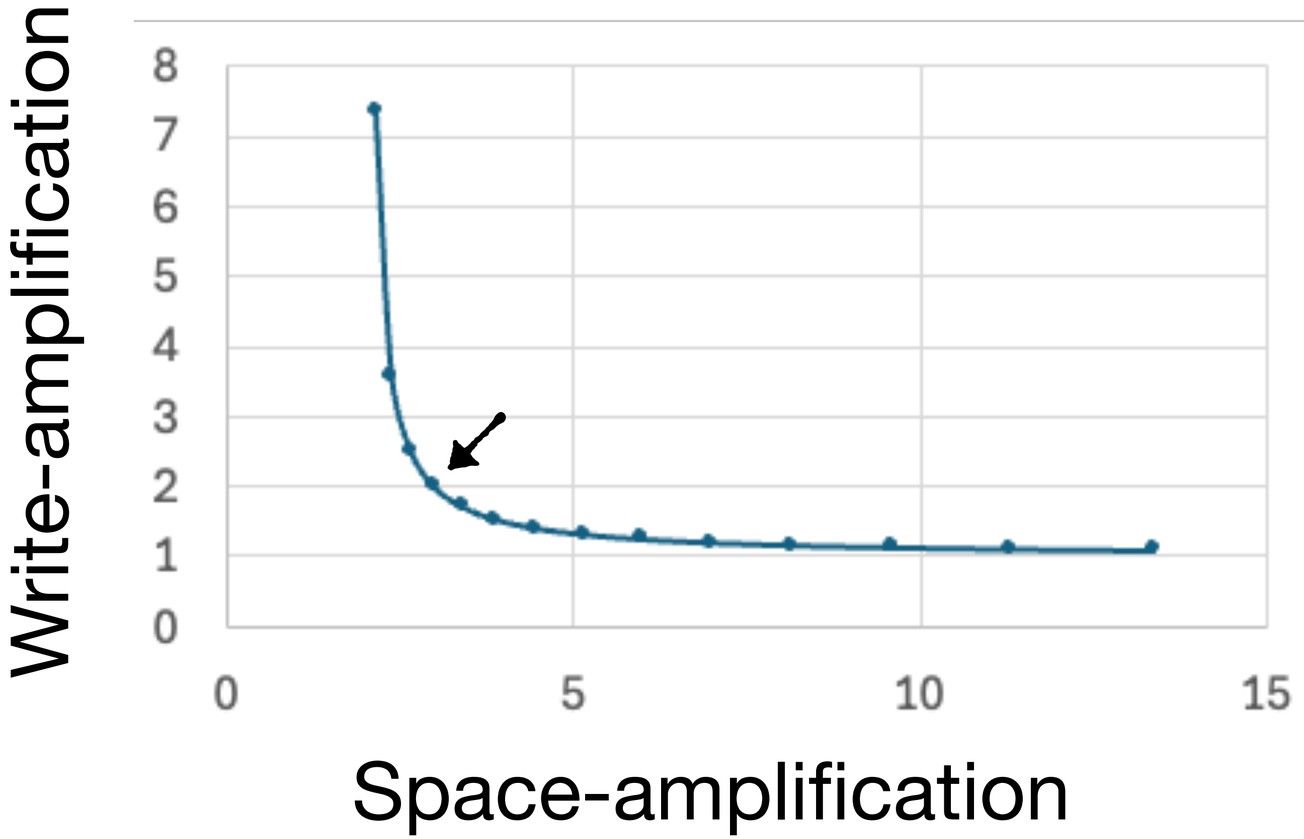

And now to new stuff

Reusing Deallocated Space

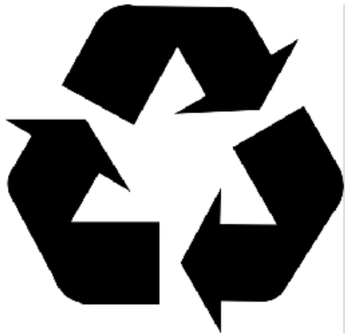

Alleviating trade-off via indirection

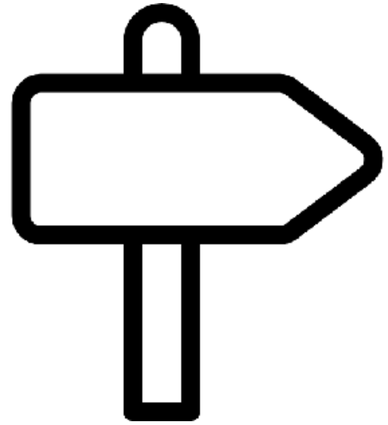

Reusing Deallocated Space

Assume G ≥ 2

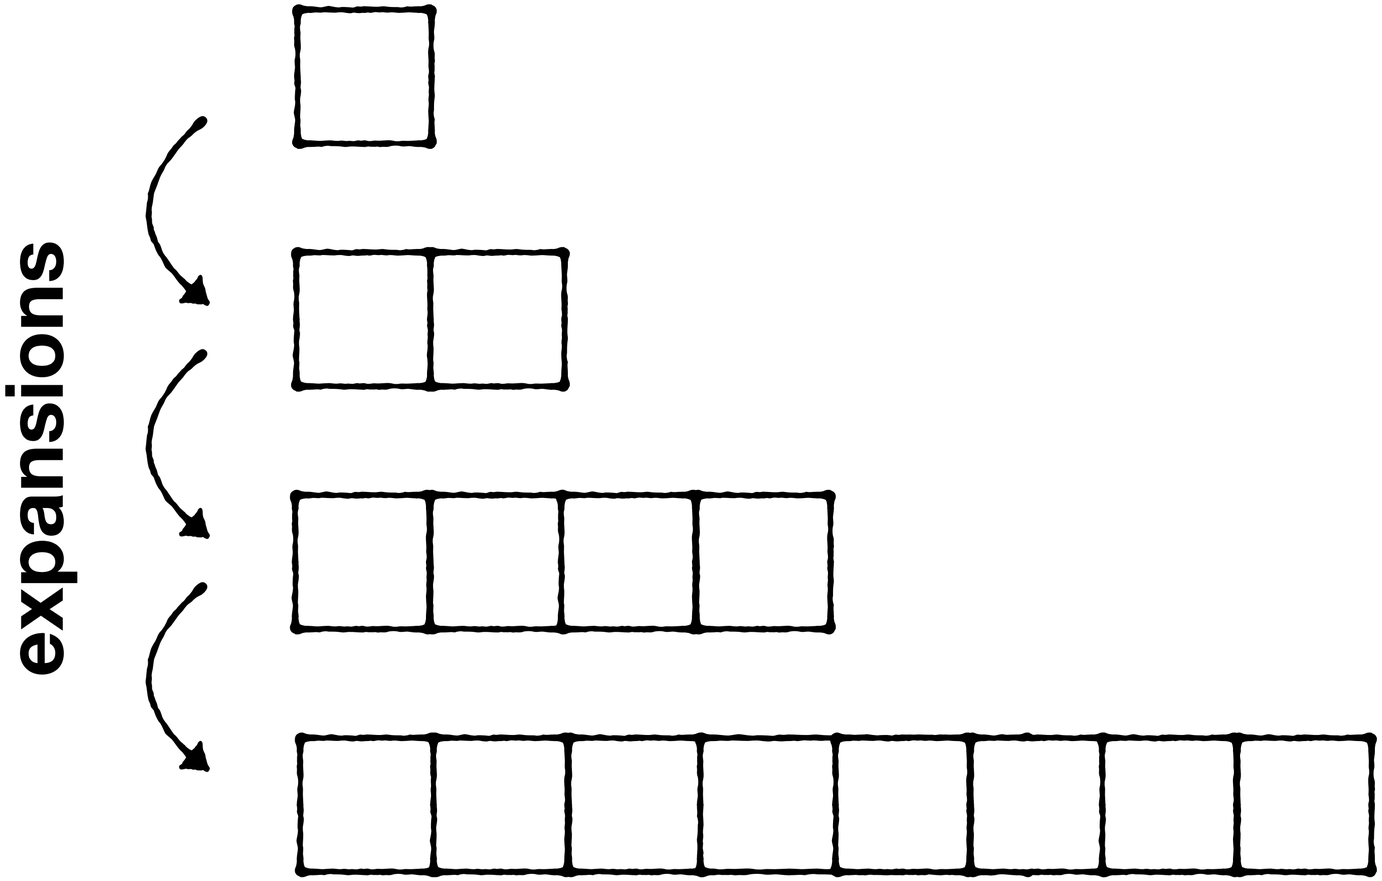

Reusing Deallocated Space (G ≥ 2)

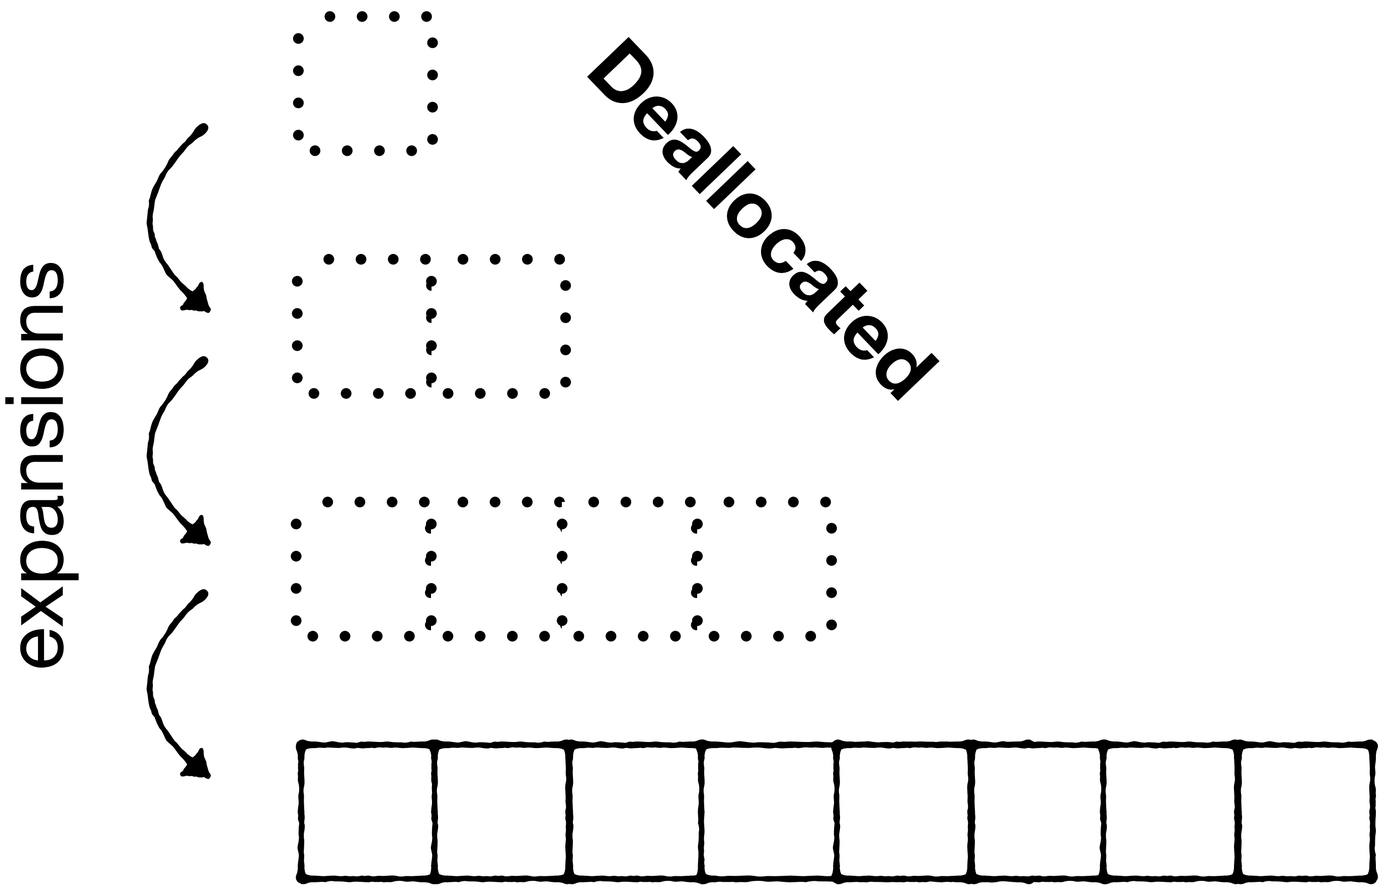

Reusing Deallocated Space

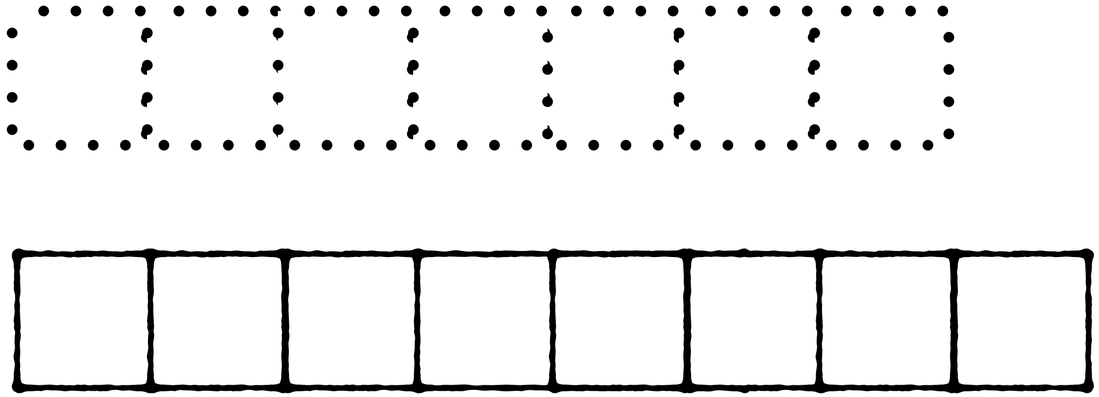

Total deallocated space

=

Total allocated space - 1

Reusing Deallocated Space

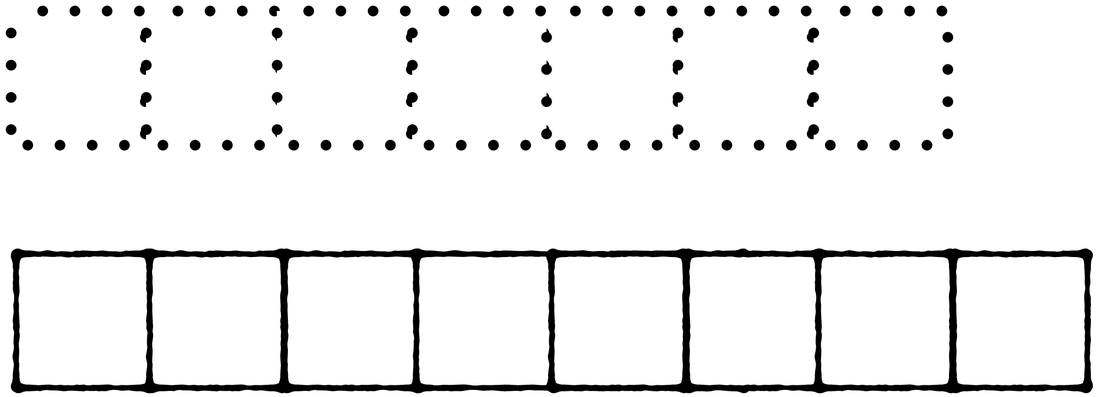

Total deallocated space

^

Total allocated space

Reusing Deallocated Space

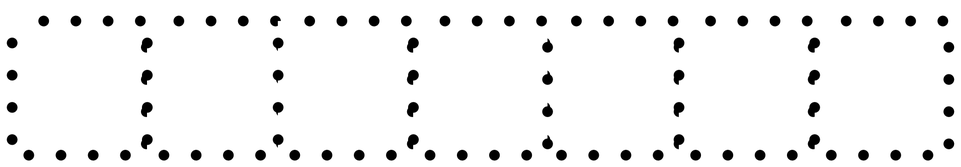

Deallocated space
Can’t be reused by new array

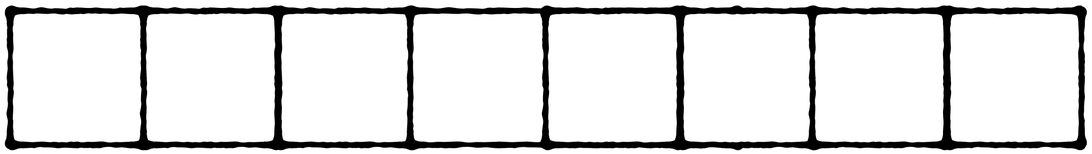

Deriving Max Space-Amp

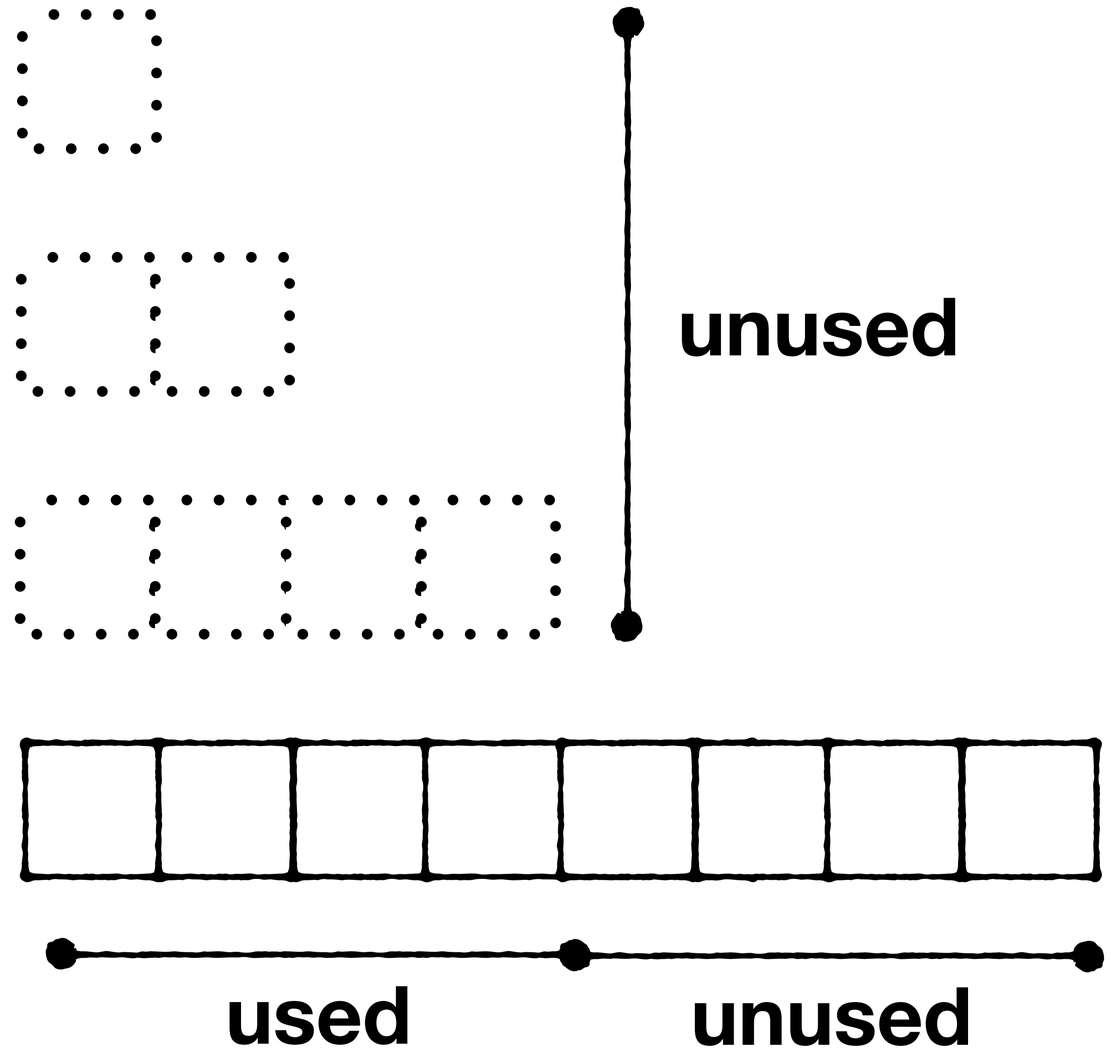

Deriving Max Space-Amp

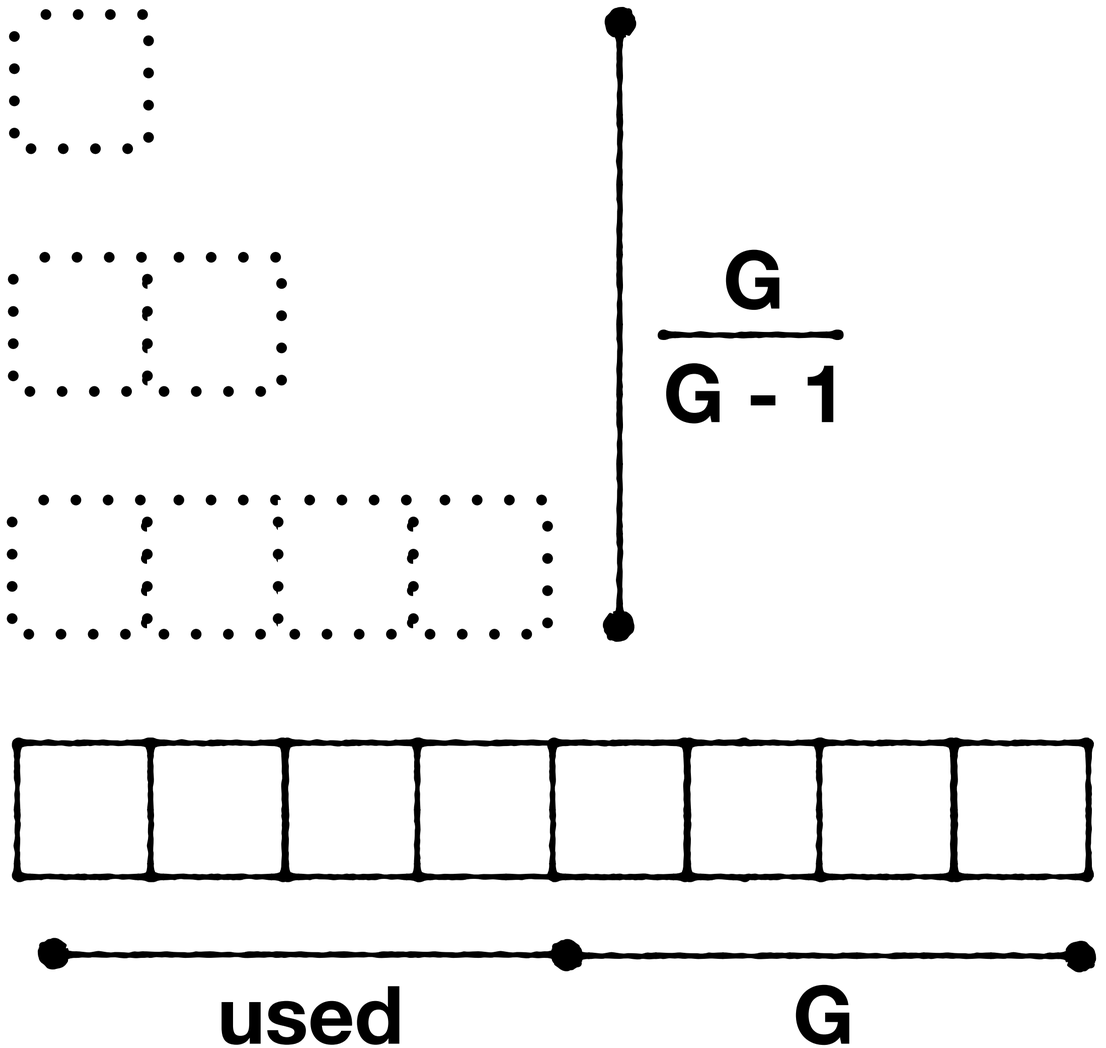

$$\text{Max Space-Amp} = \frac{\text{used + unused}}{\text{used}} = \frac{G}{G - 1} + G$$

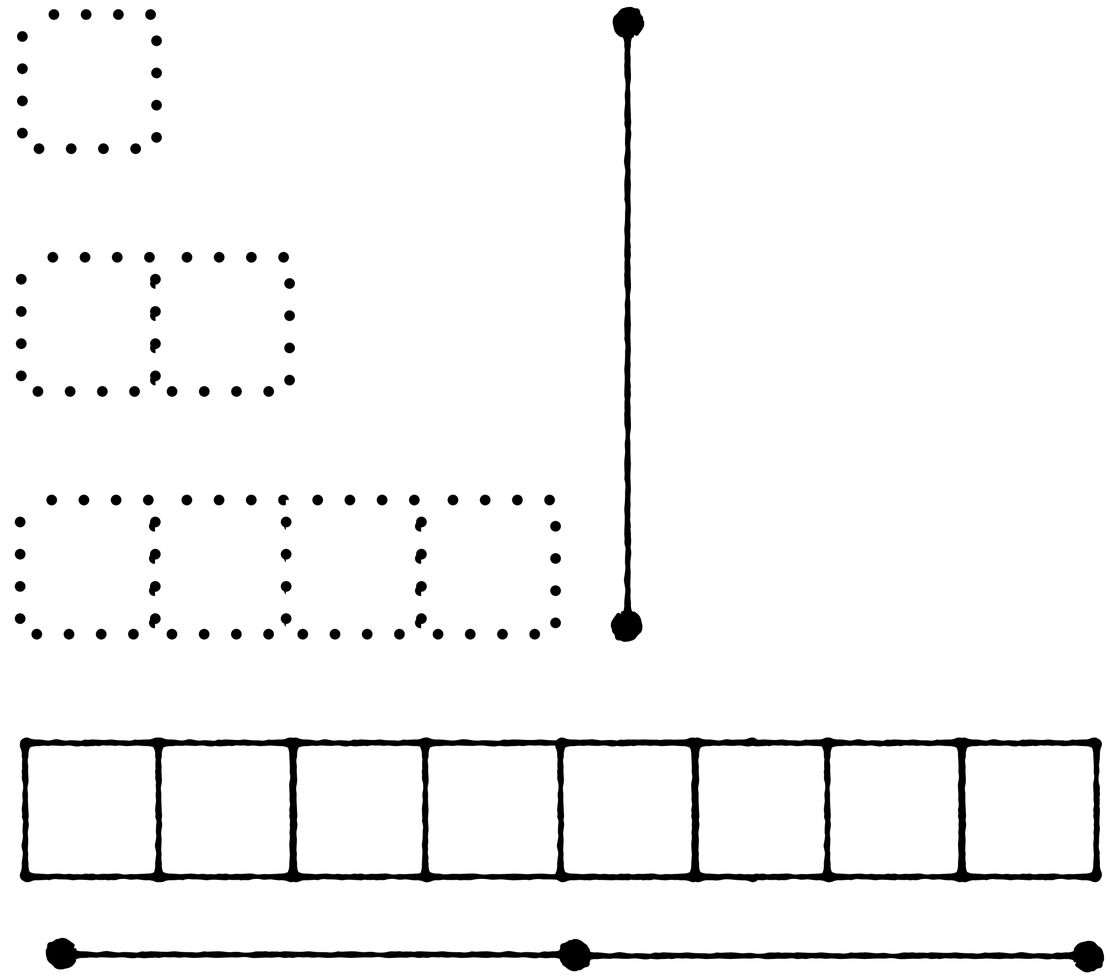

$$\text{Max Space-Amp} = \frac{\text{used + unused}}{\text{used}} = \frac{G^2}{G - 1}$$

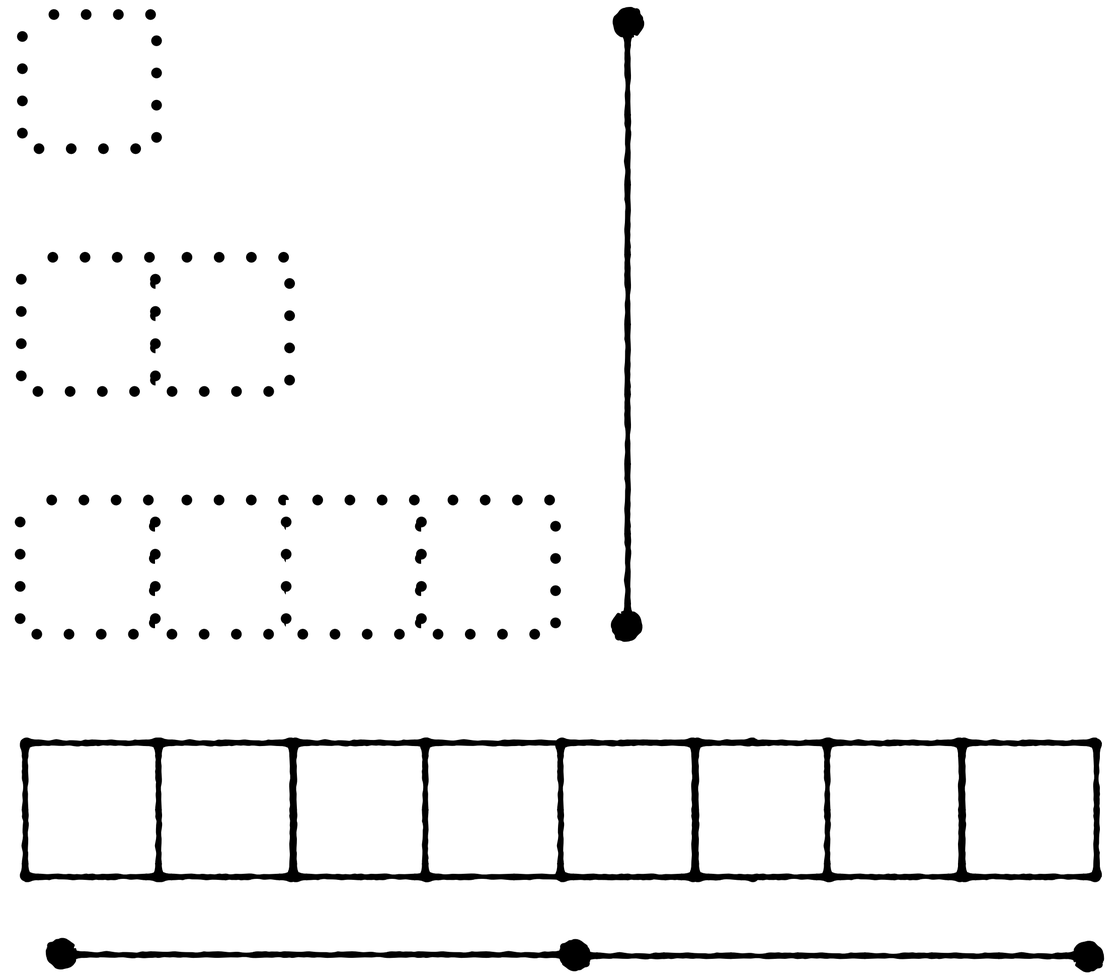

$$\text{Max Space-Amp} = \frac{\text{used + unused}}{\text{used}} = \frac{G^2}{G - 1} = 4 \quad \textbf{for G = 2}$$

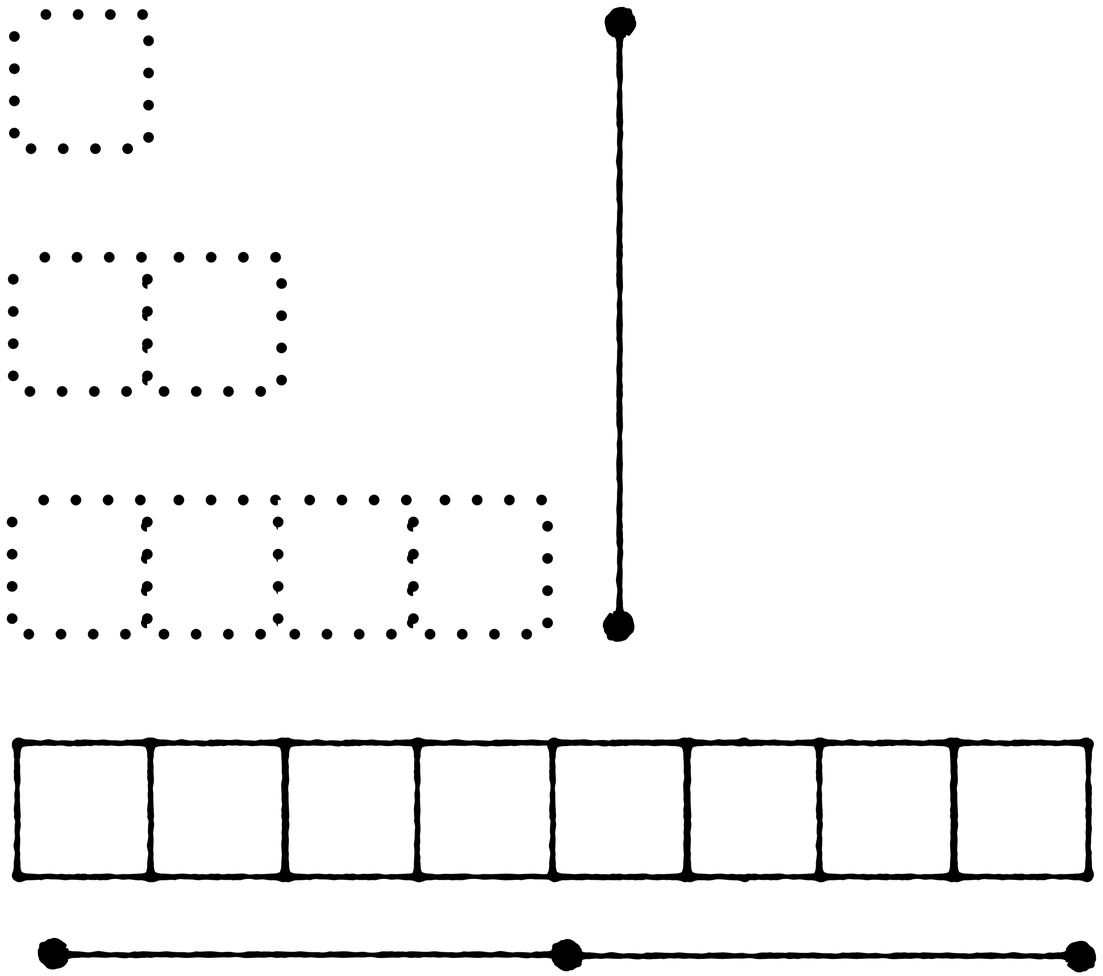

Suppose we use small size ratio,
e.g., G = 1.2

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

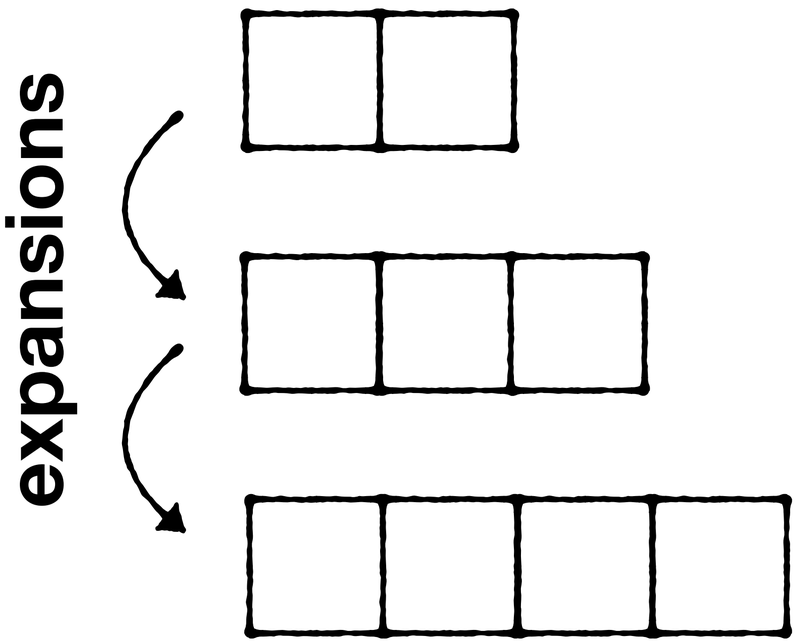

Suppose we use small size ratio, e.g., G = 1.2

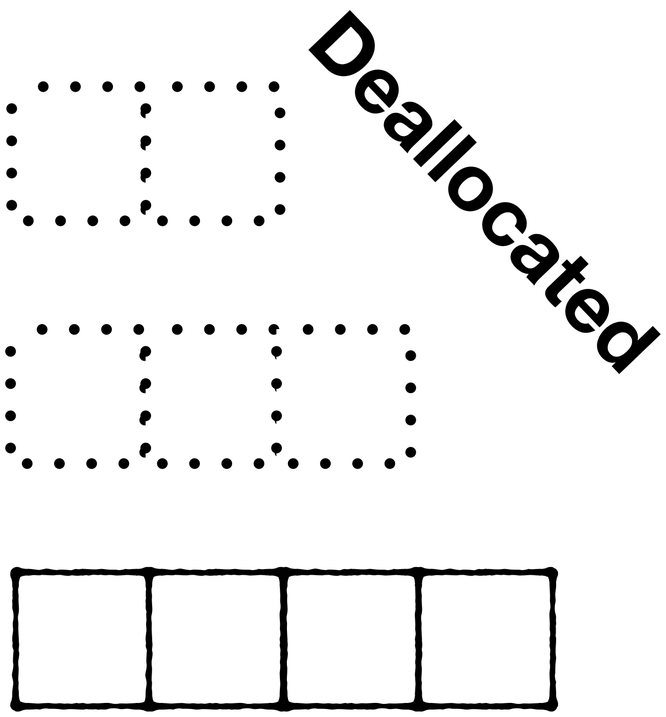

Suppose we use small size ratio,
e.g., G = 1.2

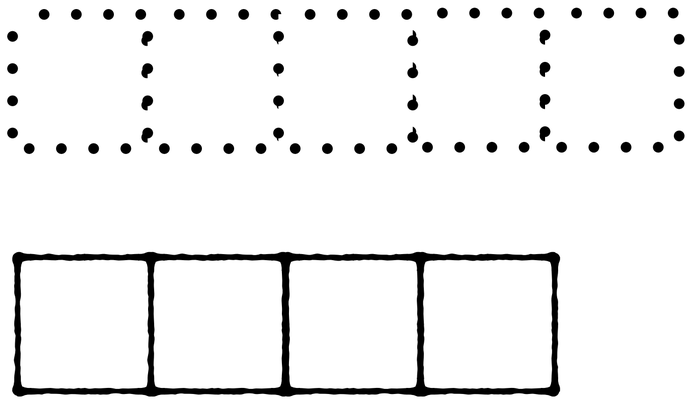

Total deallocated space

v

Total allocated space

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

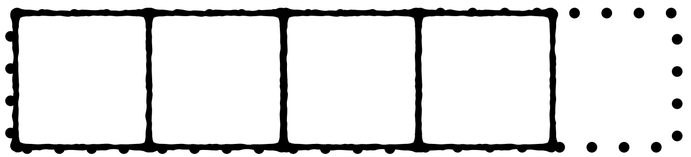

Reuse is possible

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

Wasted space

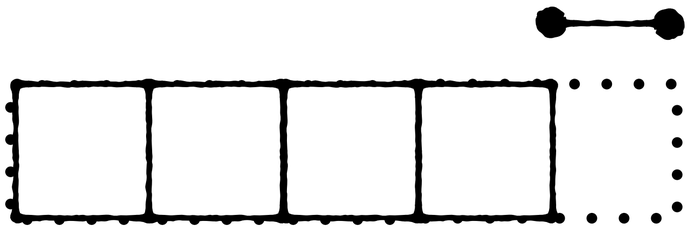

Suppose we use small size ratio,
e.g., G = 1.2

Wasted space

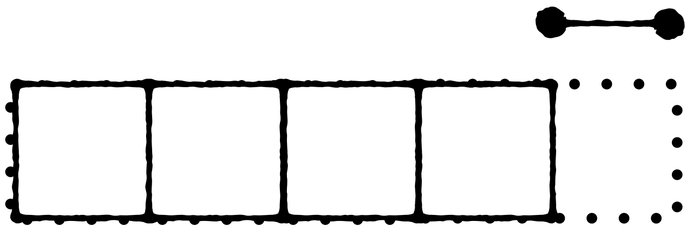

Could have expanded by larger factor

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

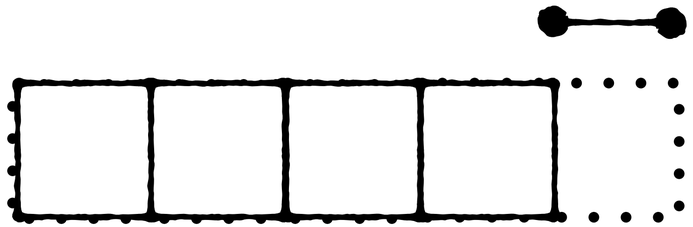

Assumptions on memory allocator

Assumptions: **Contiguous allocation**

Assumptions: **Contiguous allocation**

0 Memory addresses

Assumptions: **Contiguous allocation**

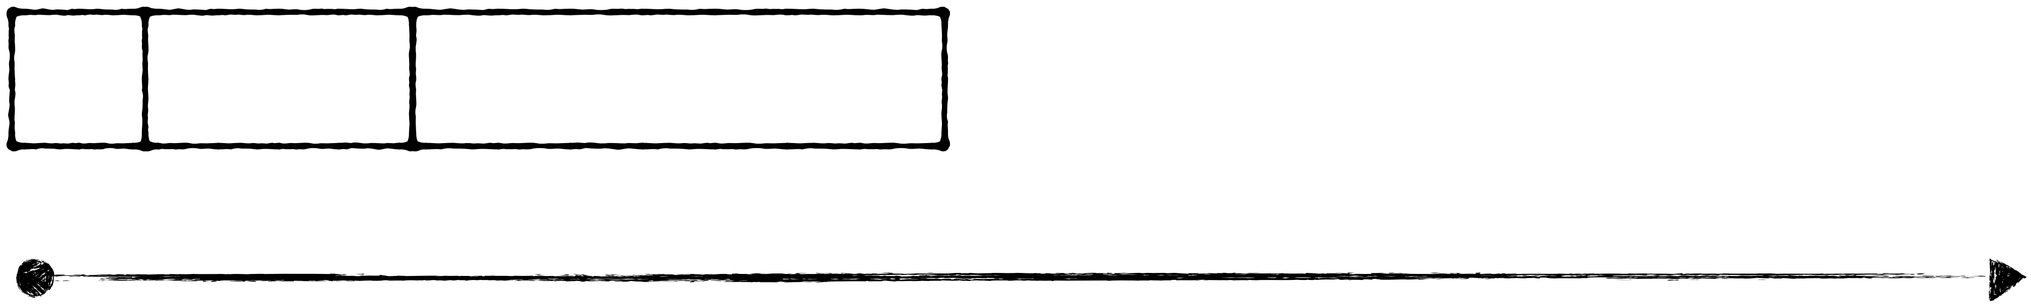

0 Memory addresses

Assumptions: **Contiguous allocation**

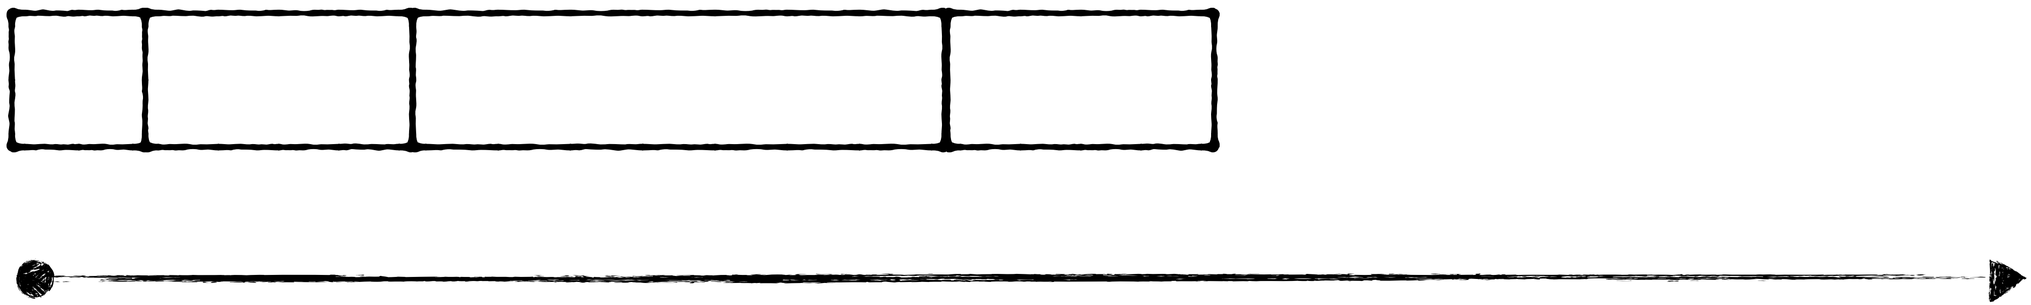

0 Memory addresses

Assumptions: **Contiguous allocation**

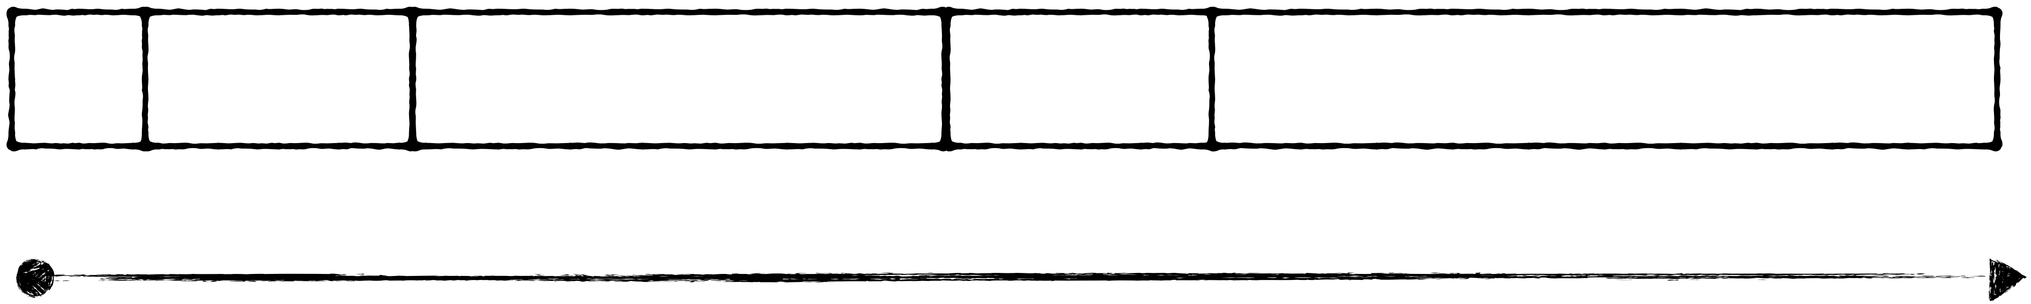

0 Memory addresses

Assumptions: Contiguous allocation

Reuse when possible

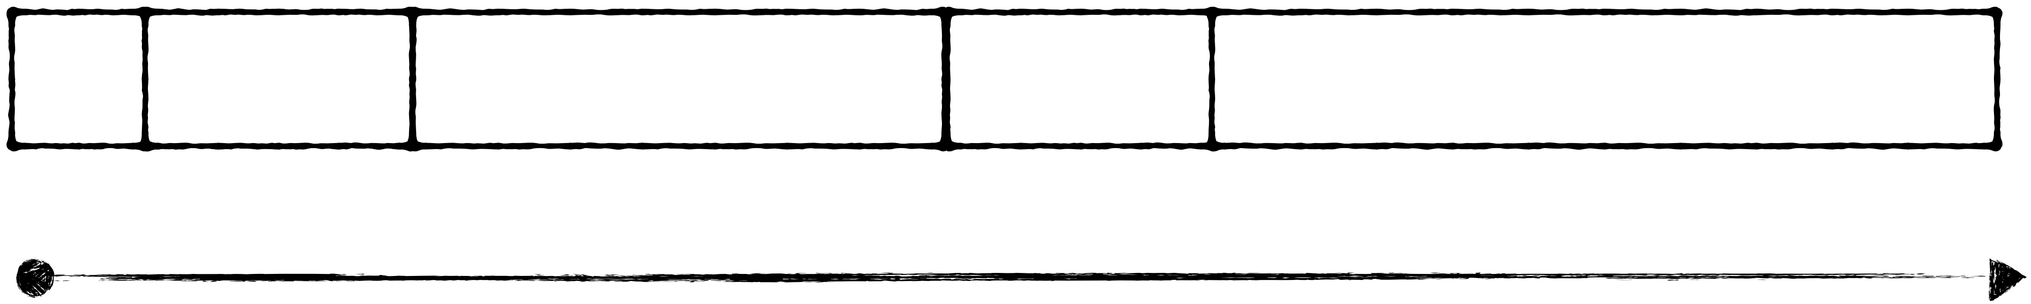

0 Memory addresses

Assumptions: Contiguous allocation

Reuse when possible

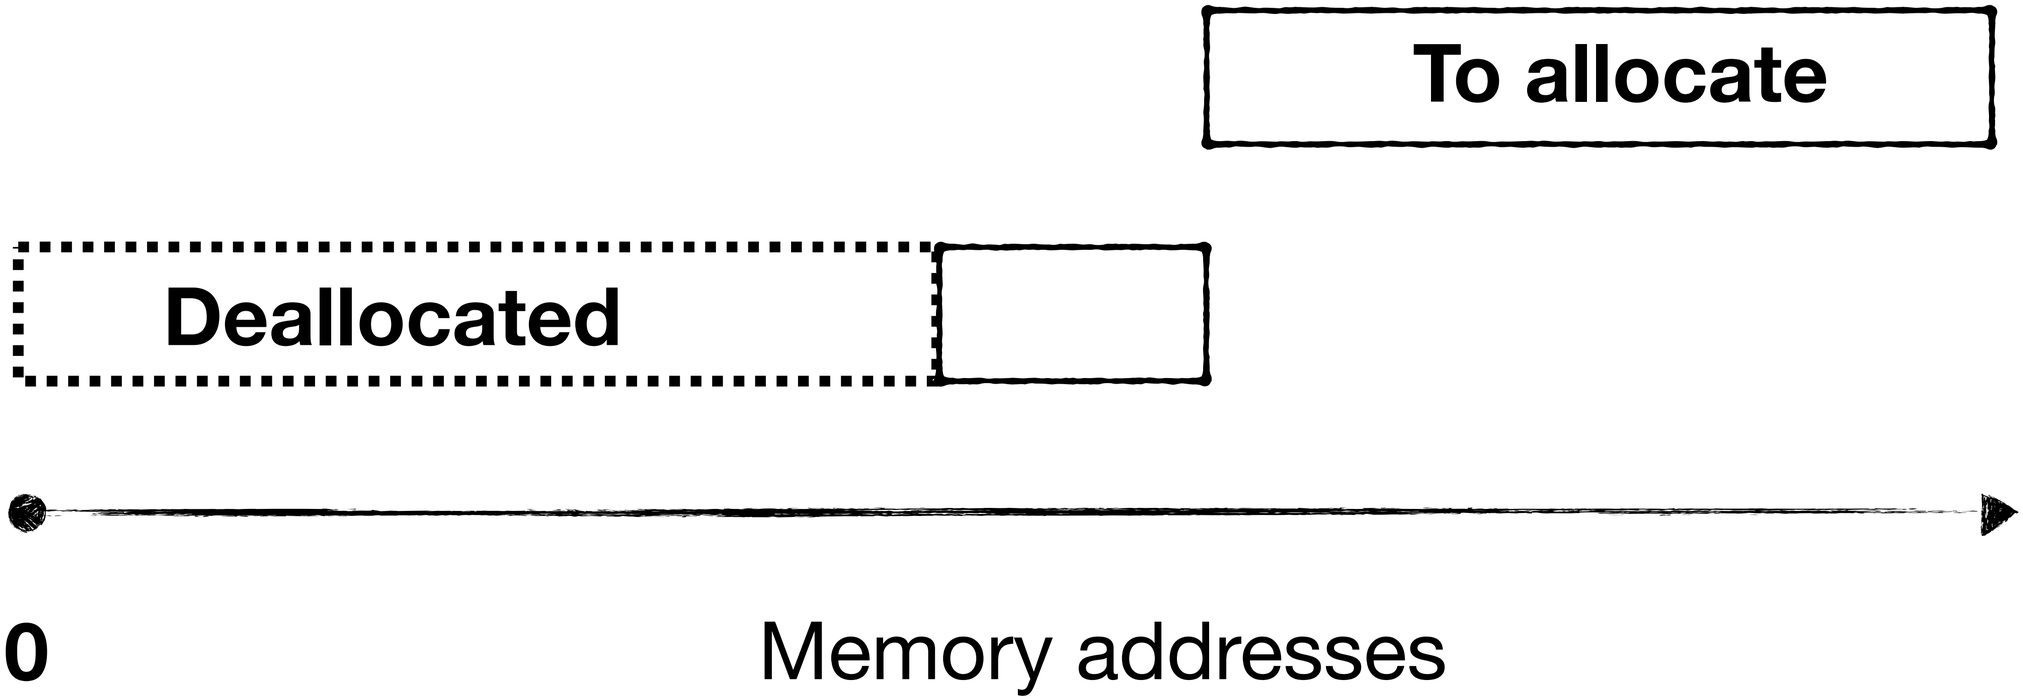

Assumptions: Contiguous allocation

Reuse when possible

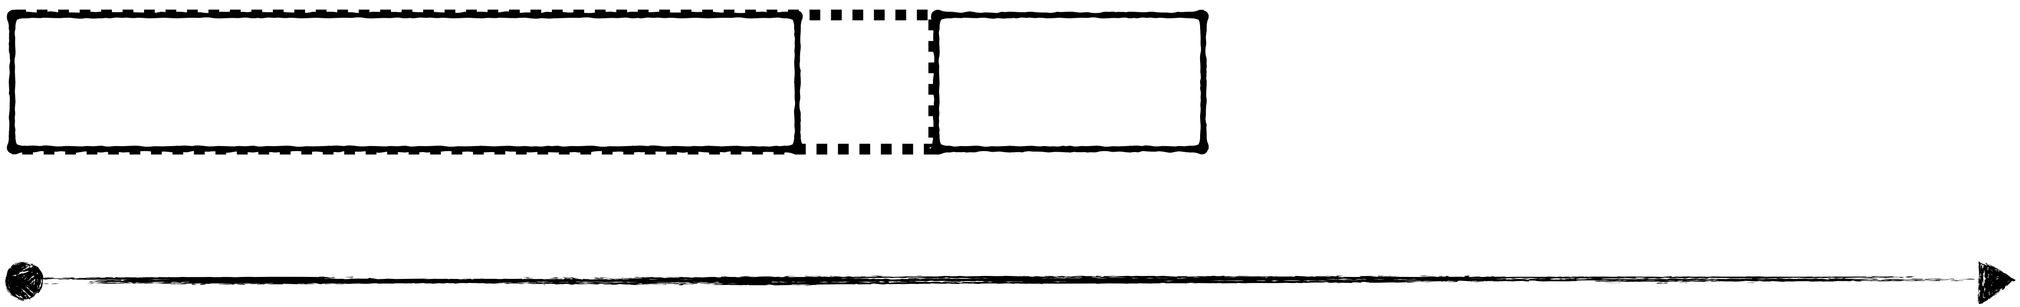

0 Memory addresses

Assumptions: Contiguous allocation

Reuse when possible

Deallocate as we copy (simplistic)

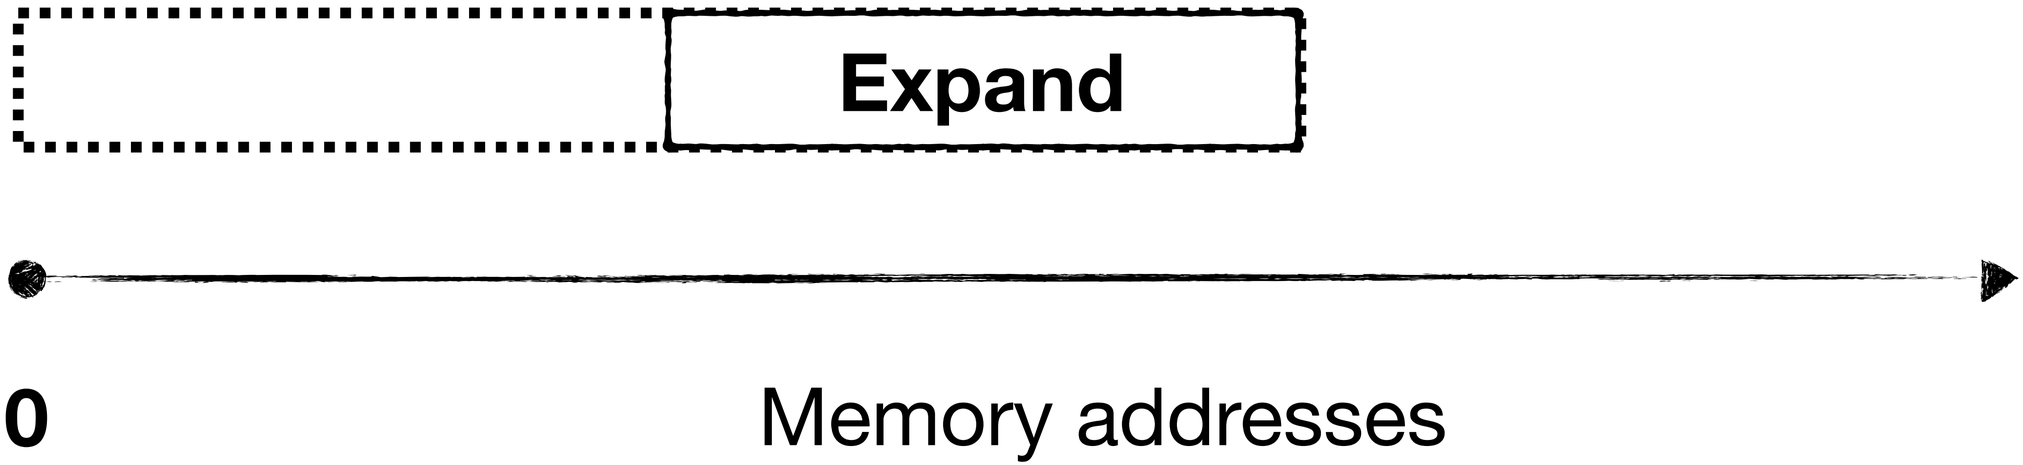

Assumptions: Contiguous allocation

Reuse when possible

Deallocate as we copy (simplistic)

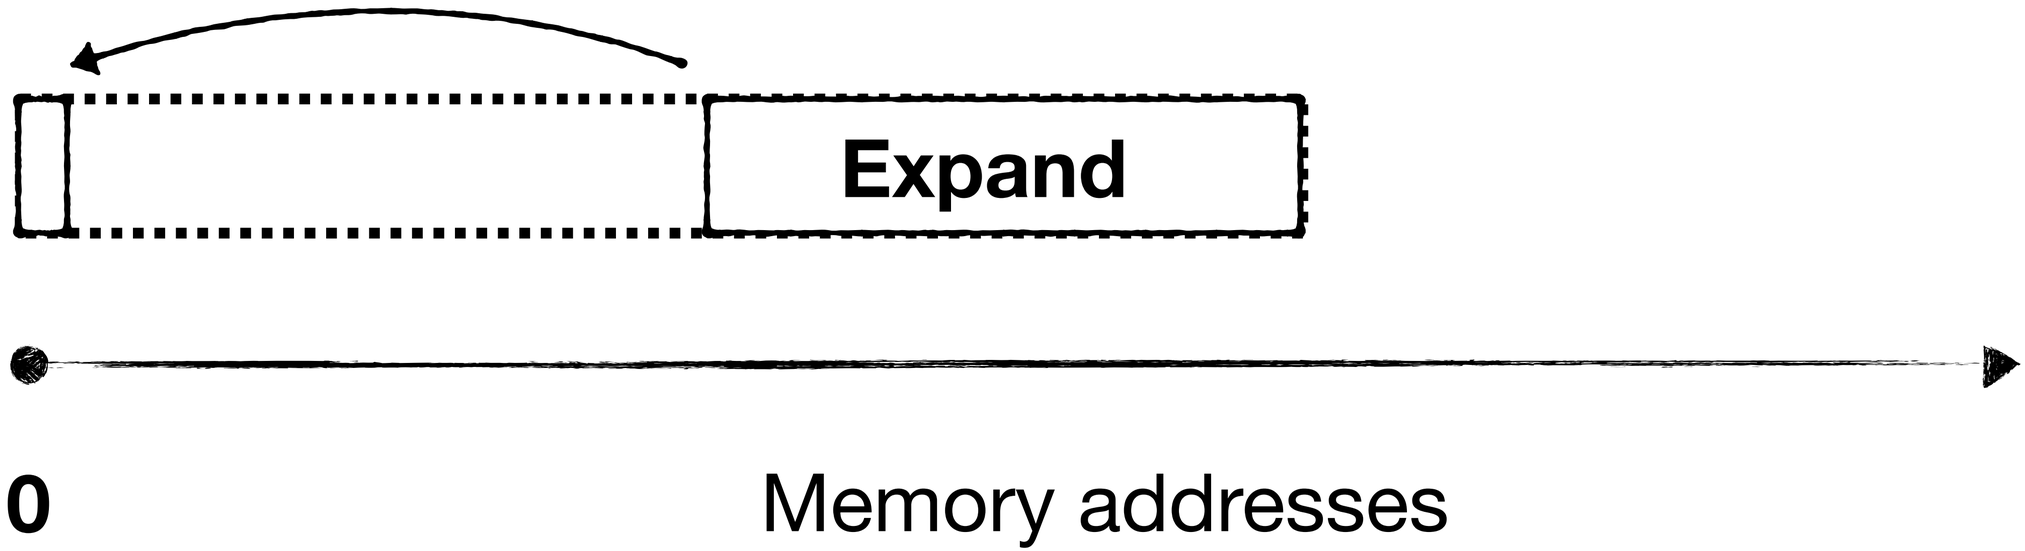

Assumptions: Contiguous allocation

Reuse when possible

Deallocate as we copy (simplistic)

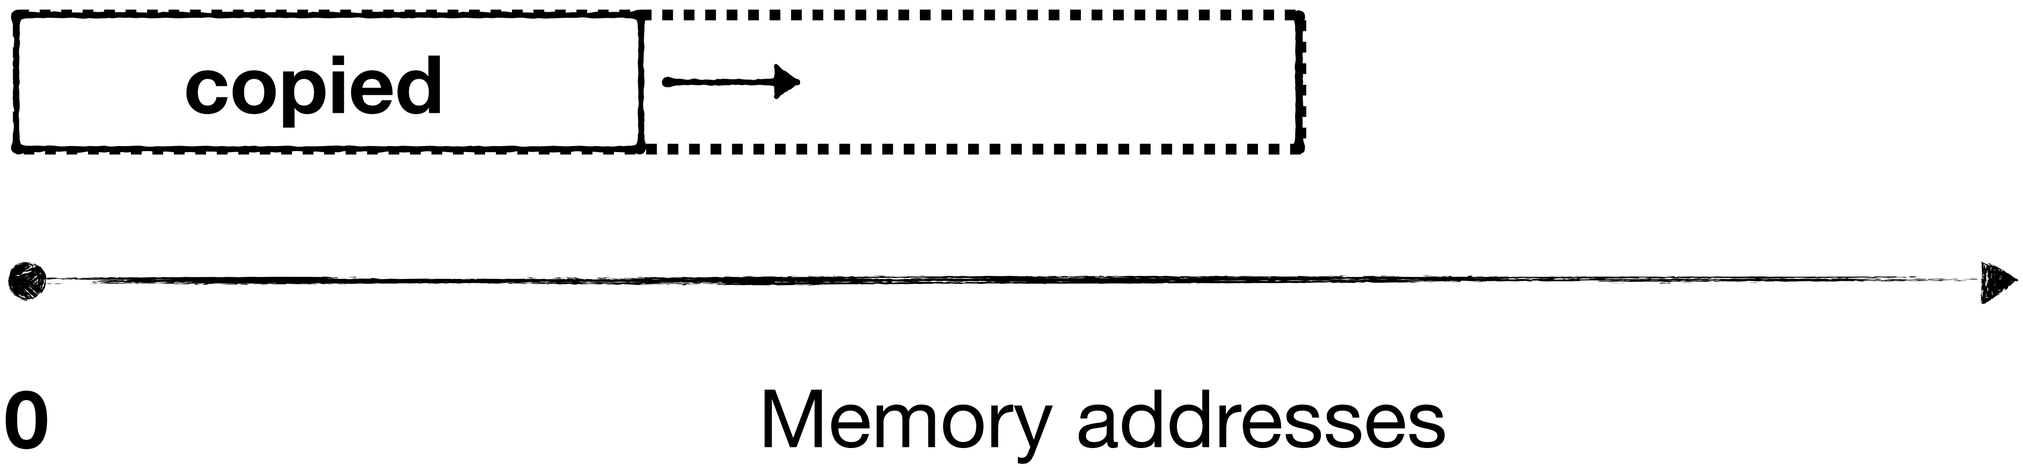

Assumptions: Contiguous allocation

Reuse when possible

Deallocate as we copy (simplistic)

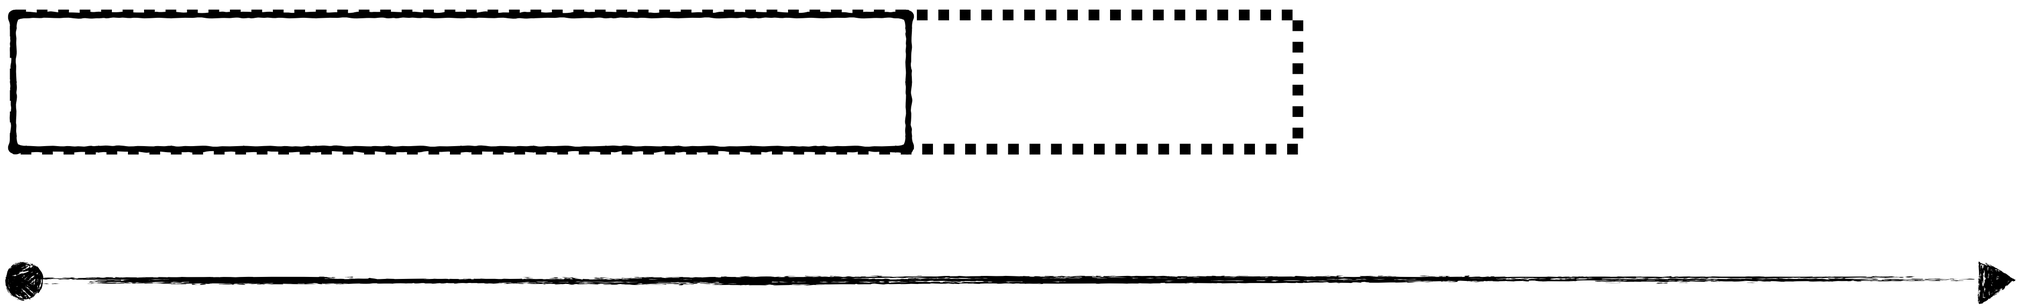

0 Memory addresses

Assumptions: Contiguous allocation

Reuse when possible

Deallocate as we copy

Can’t expand in-place

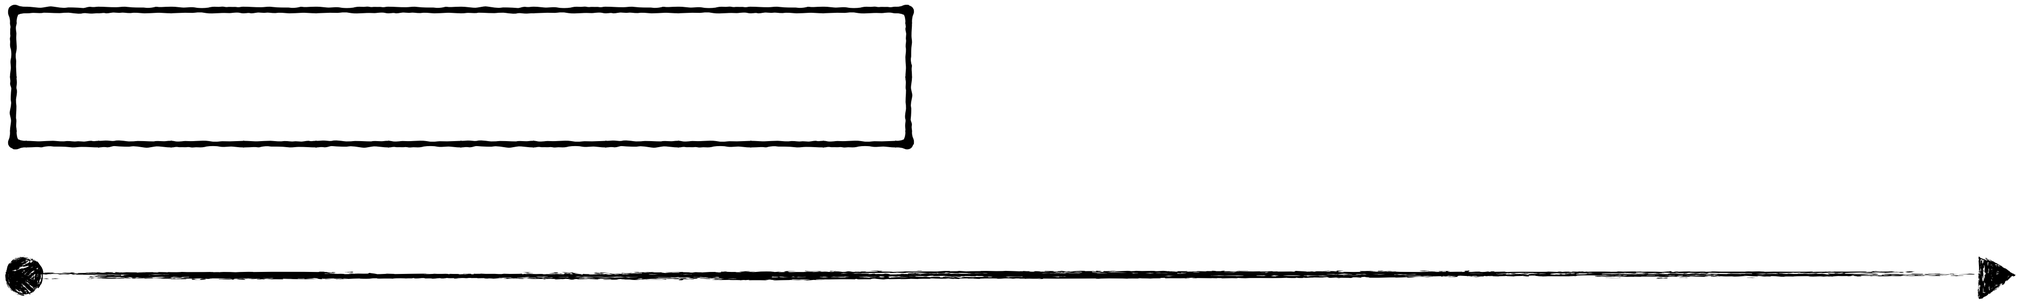

0 Memory addresses

Assumptions: Contiguous allocation

Reuse when possible

Deallocate as we copy

Can’t expand in-place

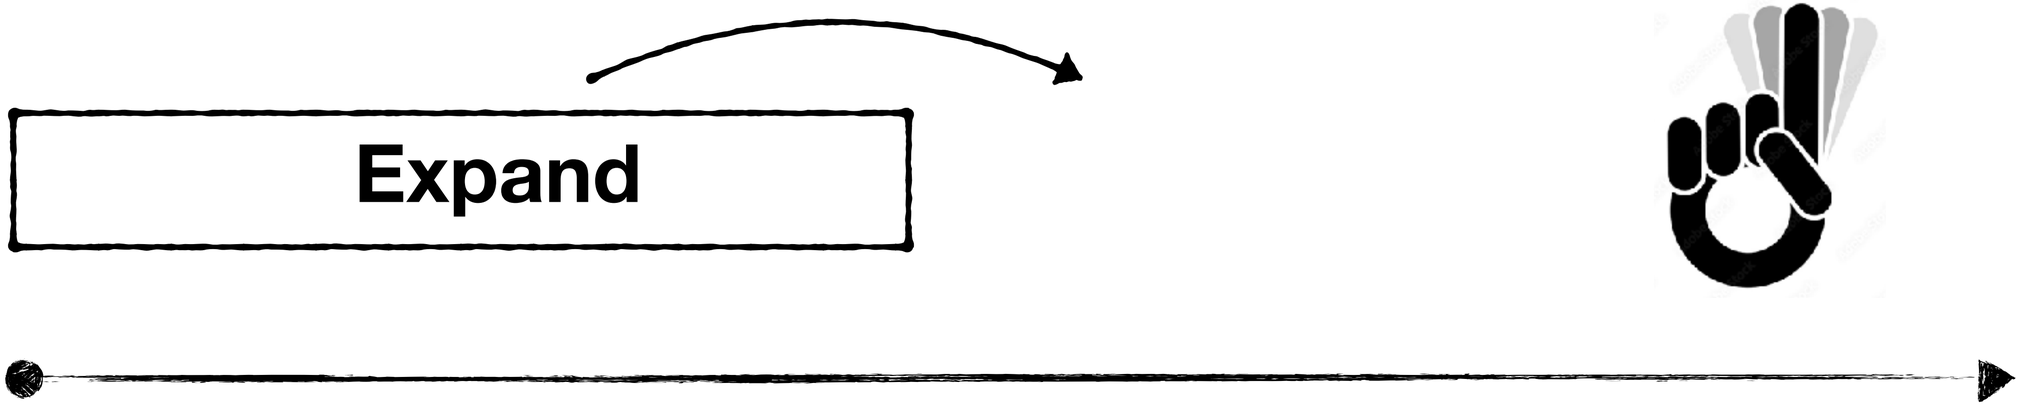

Assumptions: Contiguous allocation

Reuse when possible

Deallocate as we copy

Can’t expand in-place

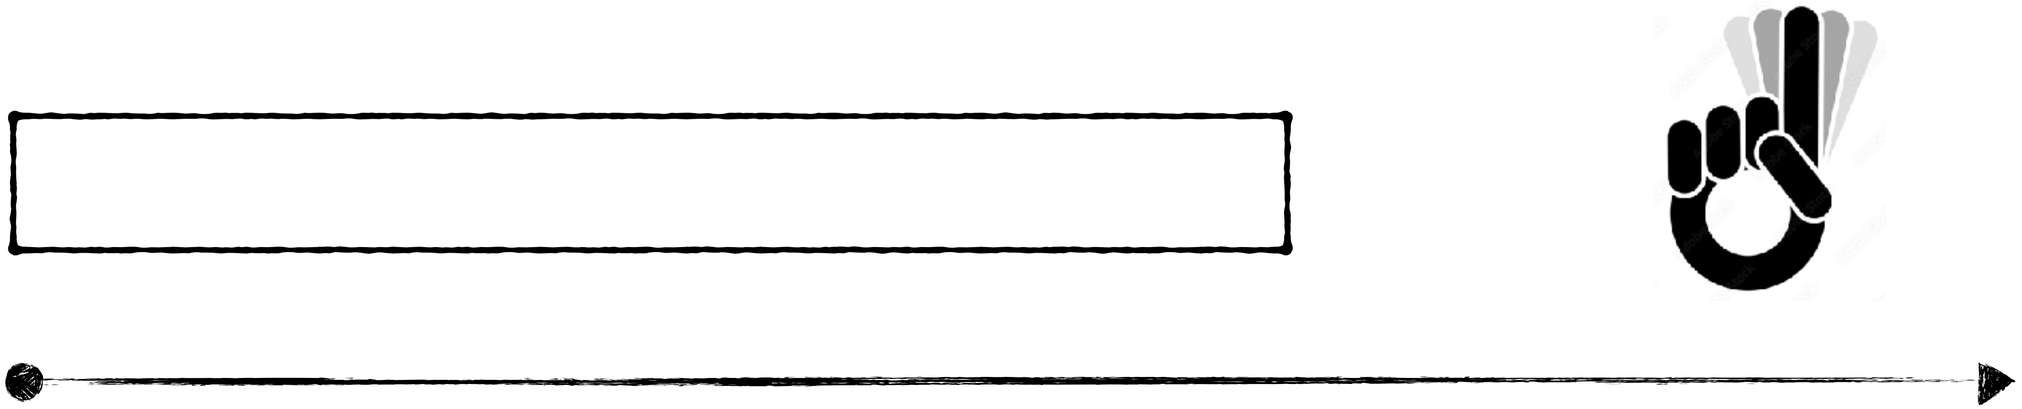

0 Memory addresses

Assumptions: Contiguous allocation

Reuse when possible

Deallocate as we copy

Can’t expand in-place

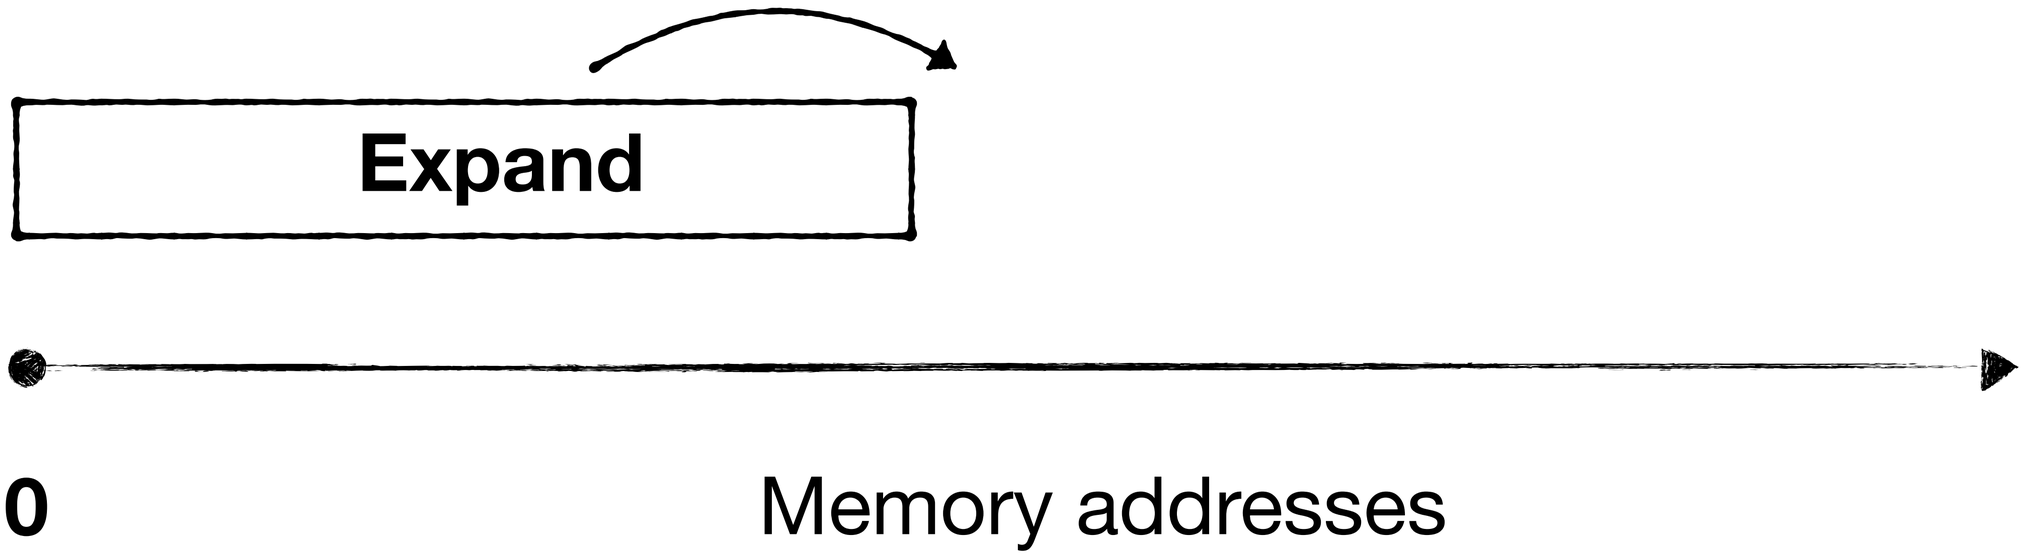

Assumptions: Contiguous allocation

Reuse when possible

Deallocate as we copy

Can’t expand in-place

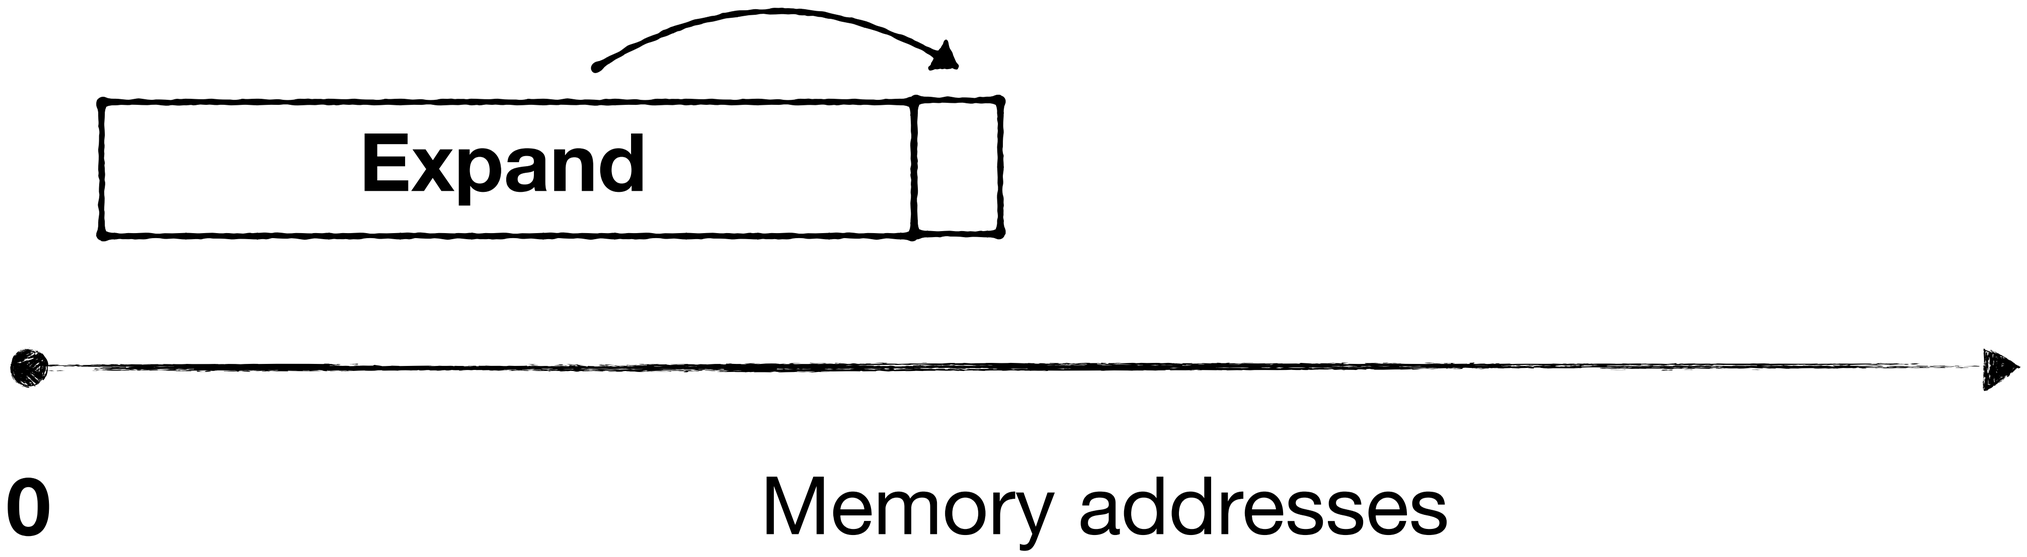

Assumptions: Contiguous allocation

Reuse when possible

Deallocate as we copy

Can’t expand in-place

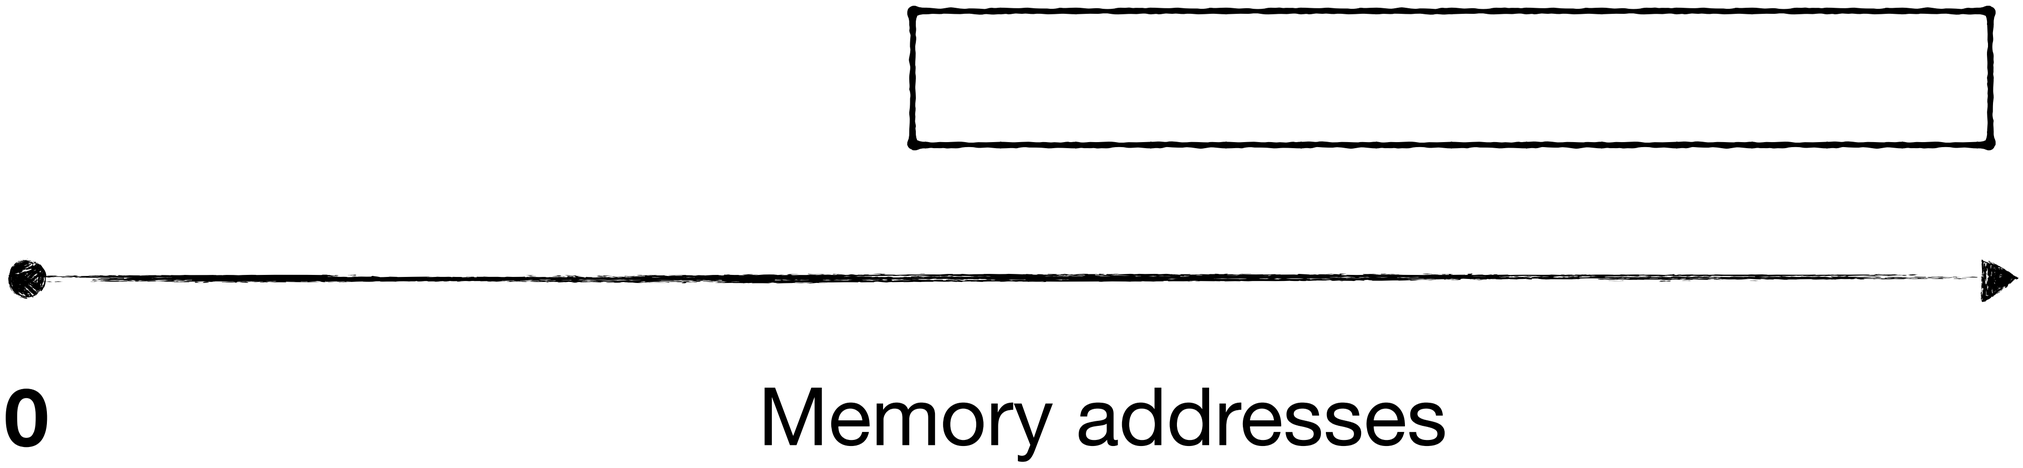

For which growth factor, do we perfectly reuse the space?

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

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

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

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

Subject to:

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

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: → 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

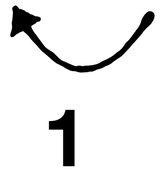

$$\text{Size}_{i-2} + \text{Size}_{i-1} = \text{Size}_{i}$$

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

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

1.615

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

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

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” ϕ = **1.618**033988749....

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

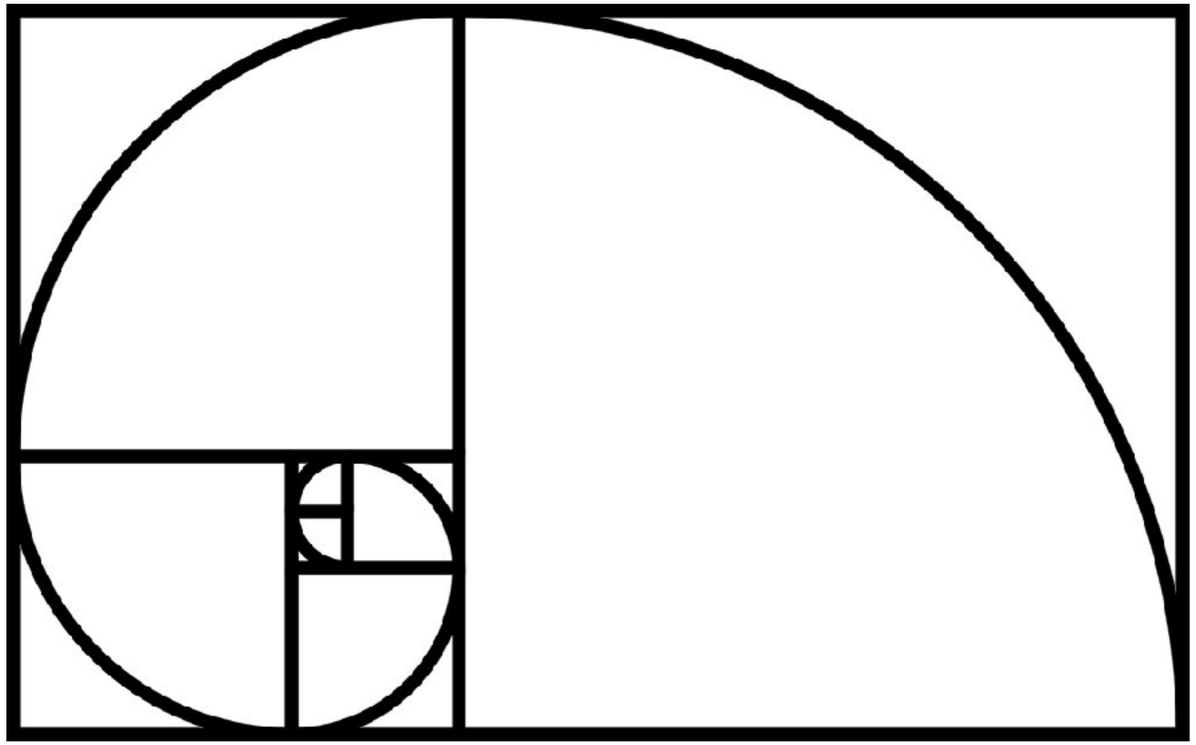

Ratio converges to the “Golden Ratio” $\phi = \frac{1+\sqrt{5}}{2}$

Golden Spiral

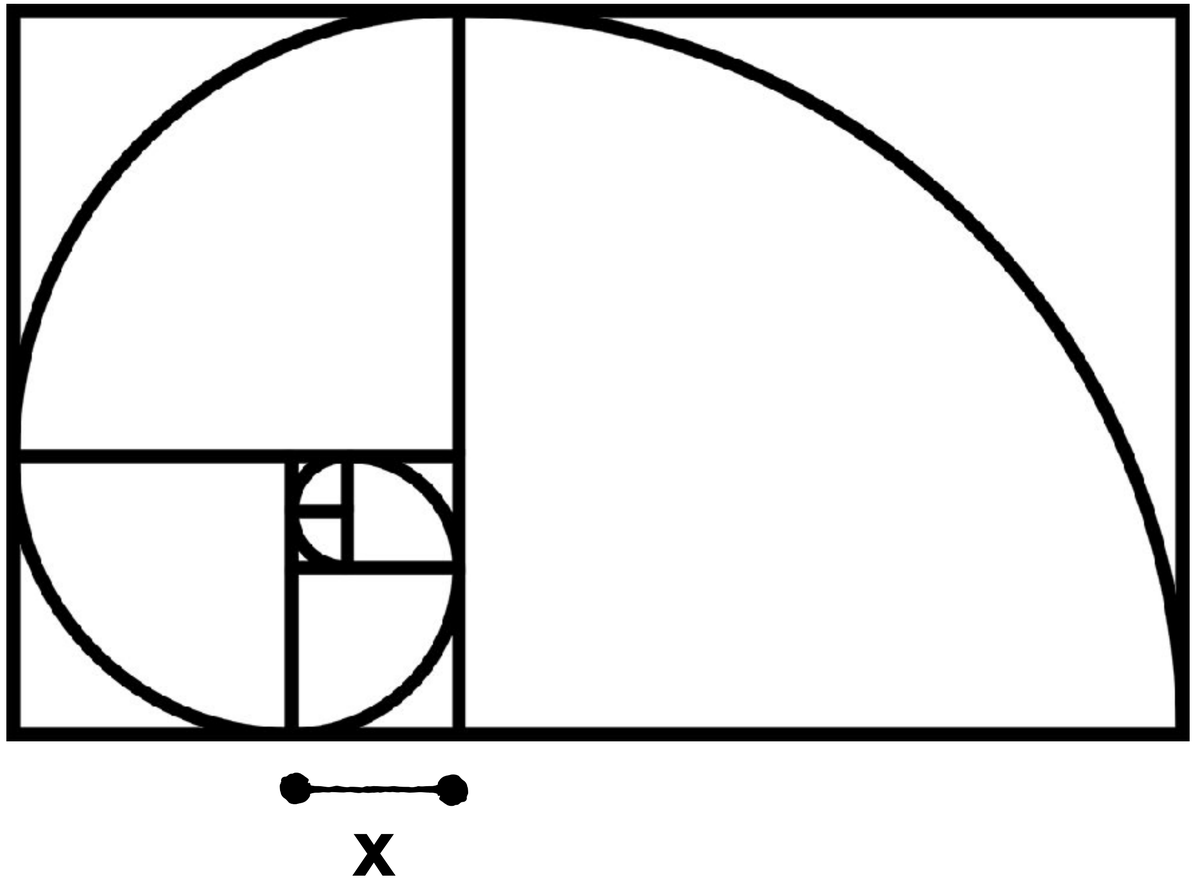

Ratio converges to the “Golden Ratio” $\phi = \frac{1+\sqrt{5}}{2}$

Golden Spiral

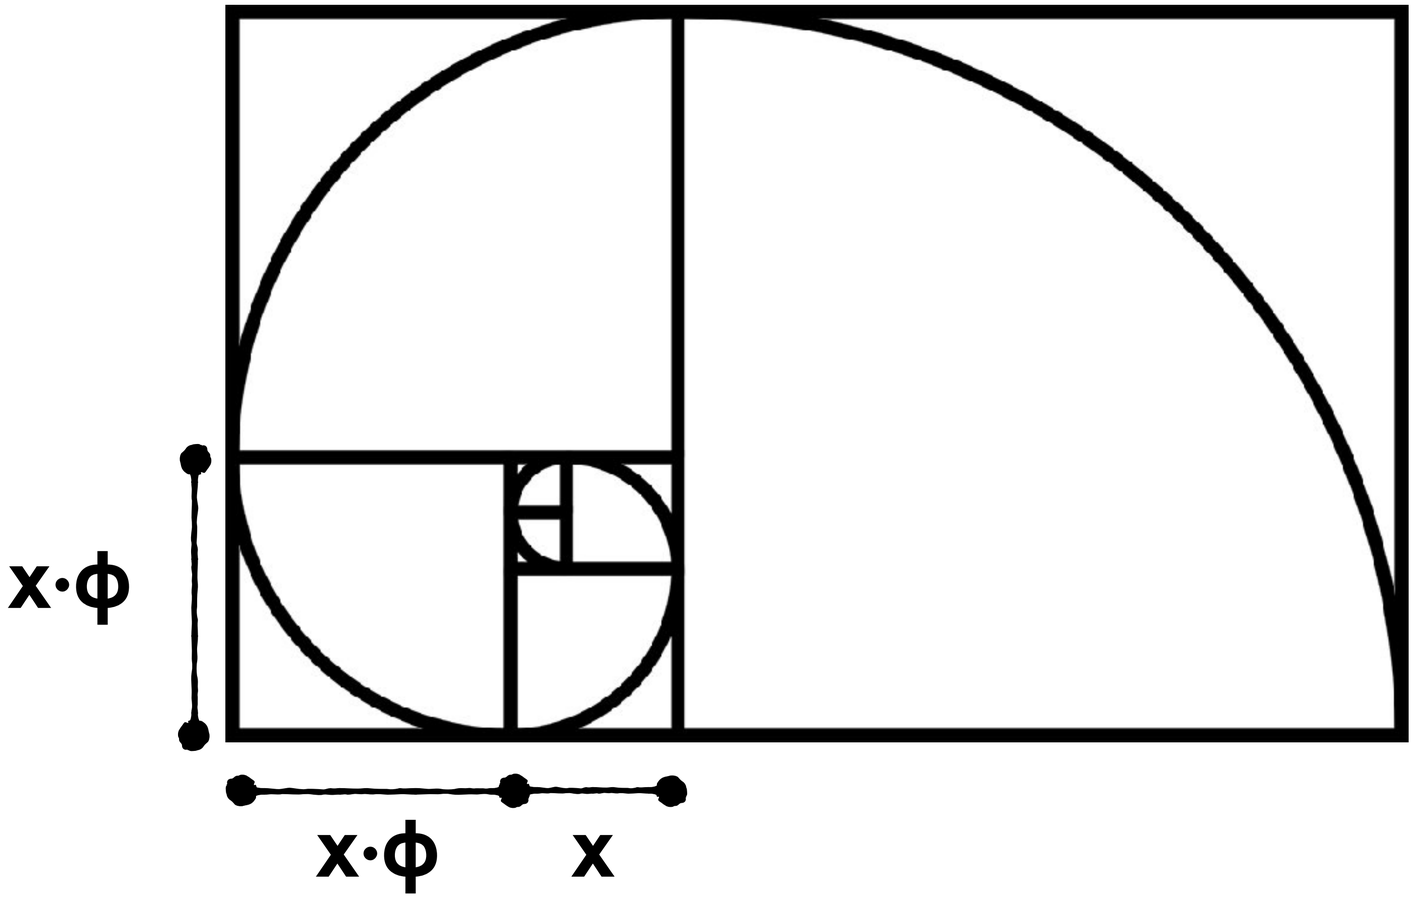

Golden Spiral

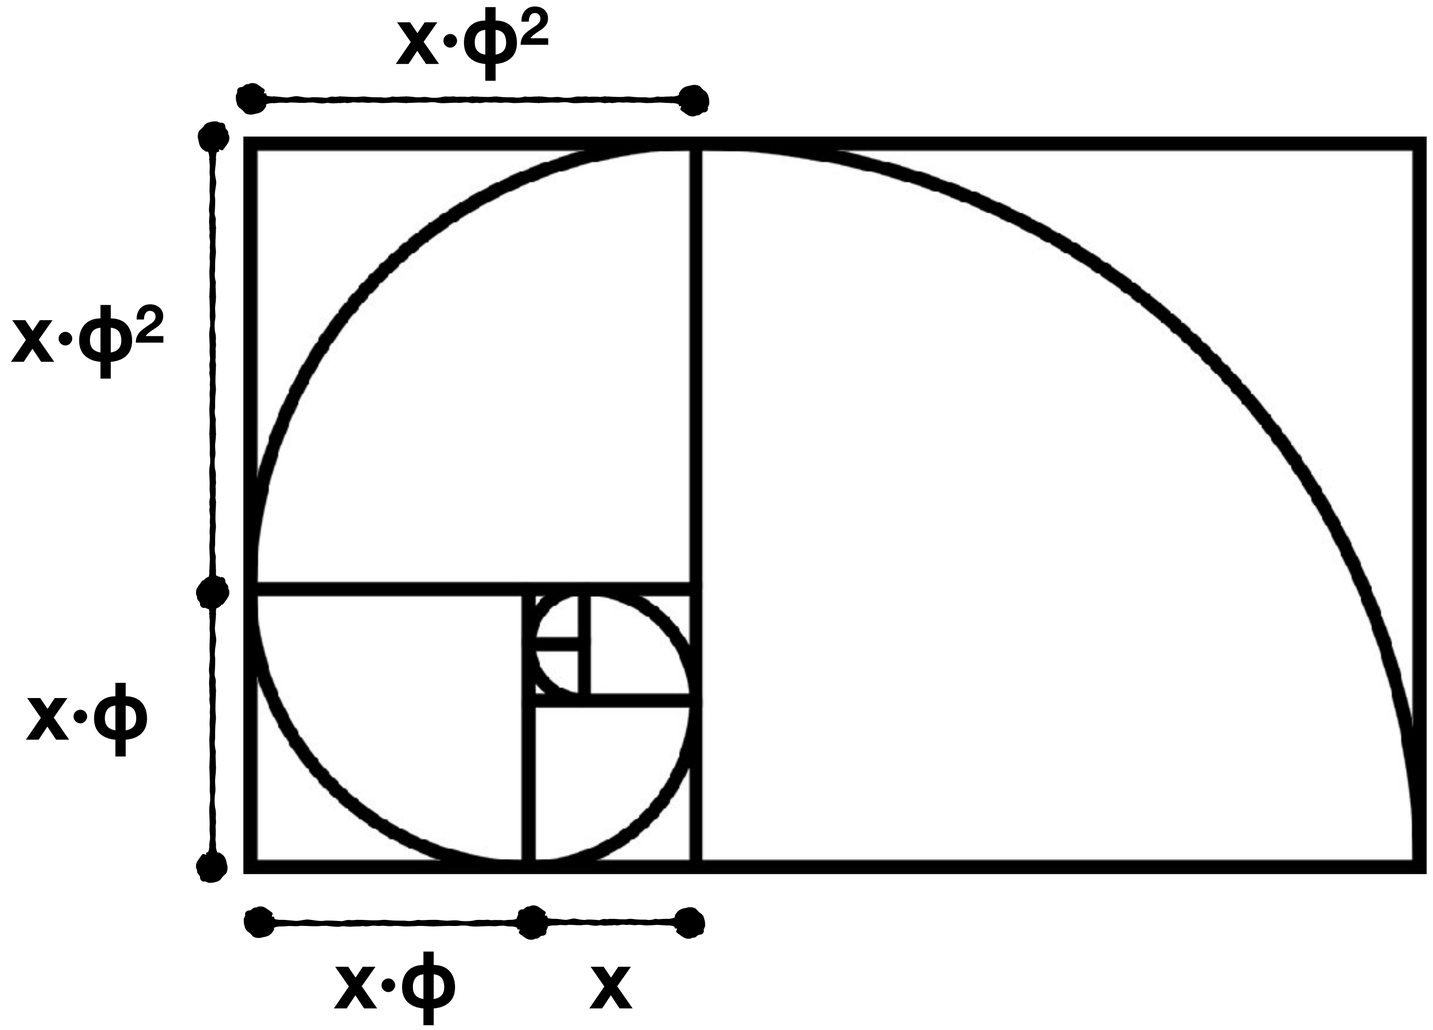

Golden Spiral

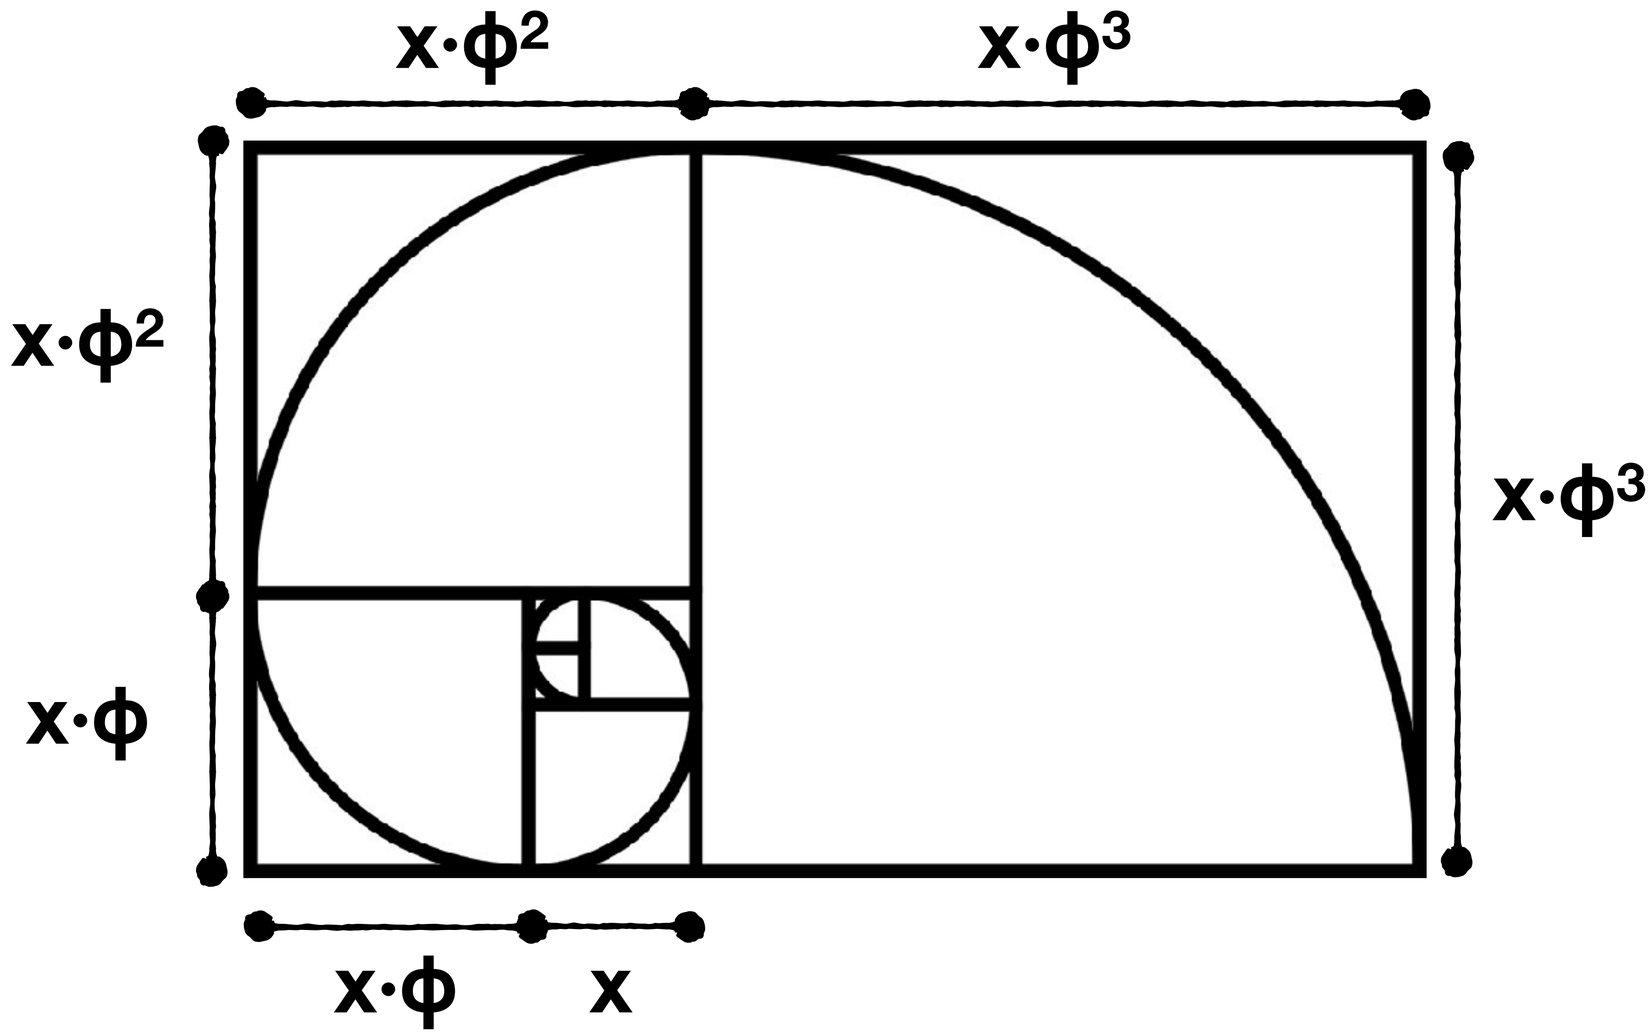

Golden Spiral

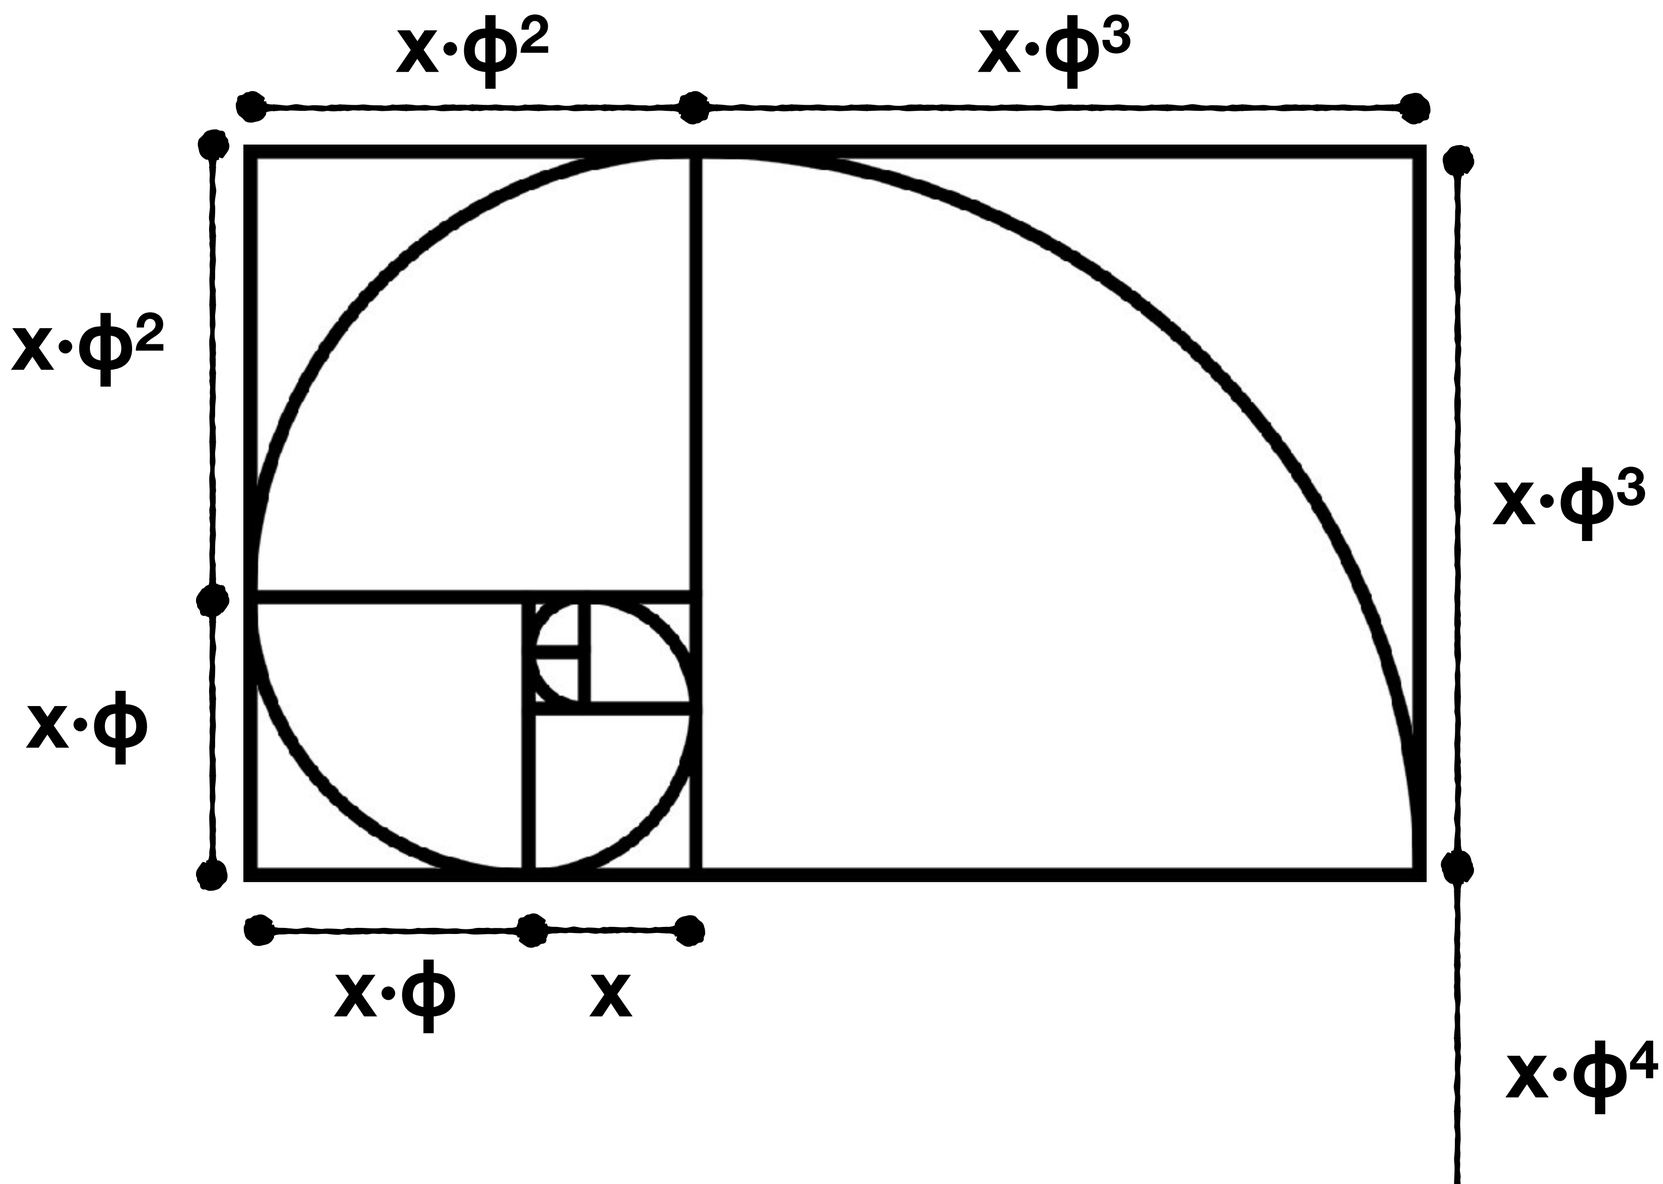

Golden Spiral

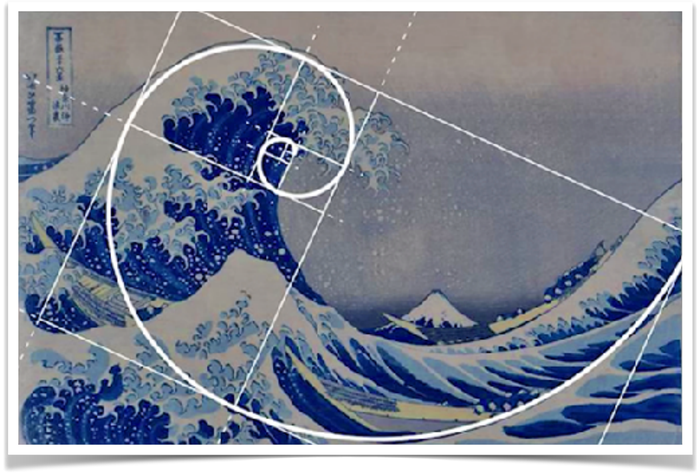

Art

The Great Wave off Kanagawa

Golden Spiral

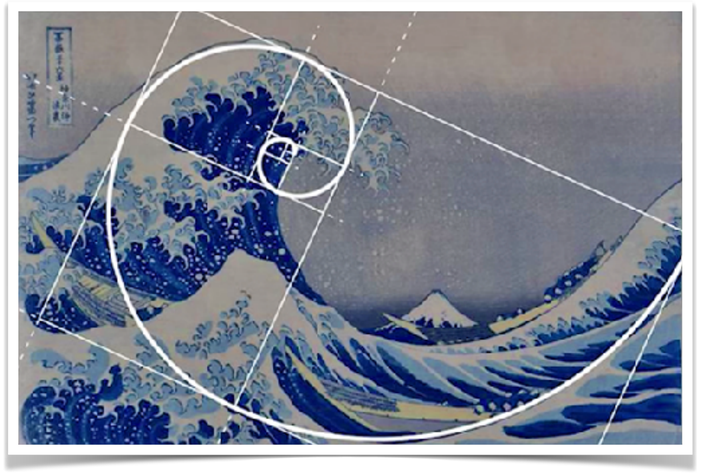

Art

The Great Wave
off Kanagawa

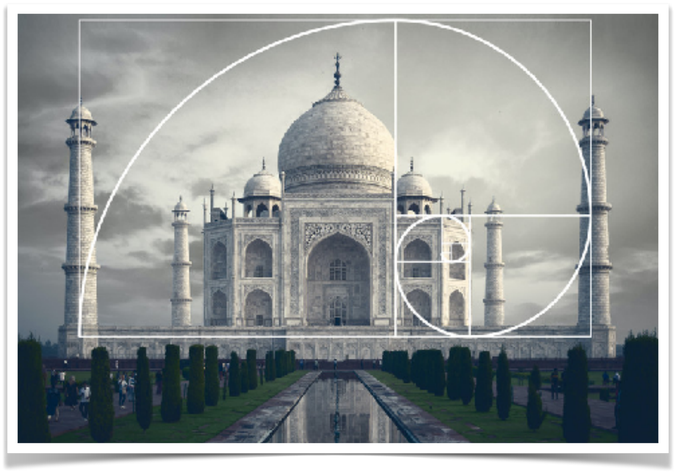

Architecture

Taj Mahal

Golden Spiral

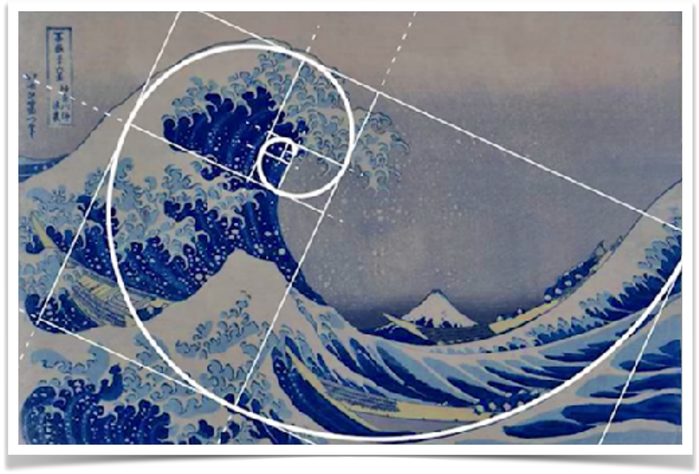

Art

The Great Wave
off Kanagawa

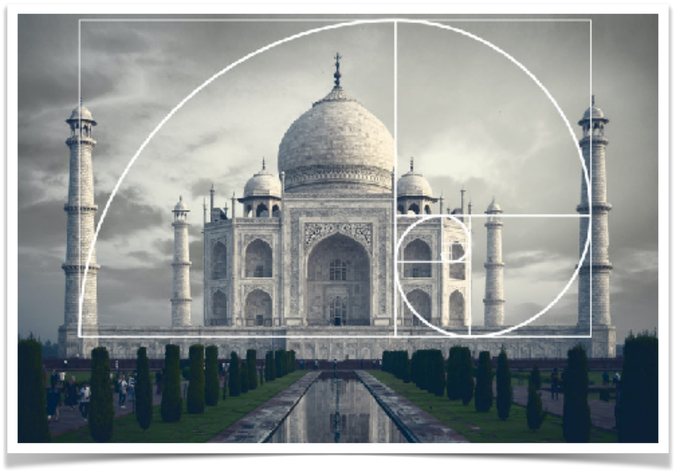

Architecture

Taj Mahal

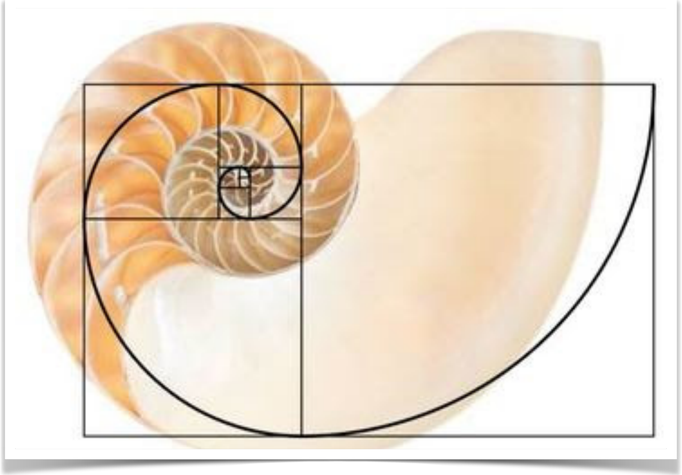

Nature

Nautilus Shell

Golden Spiral

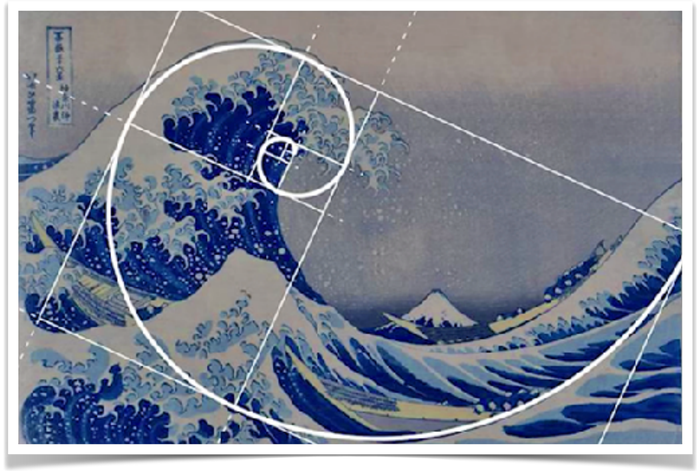

Art

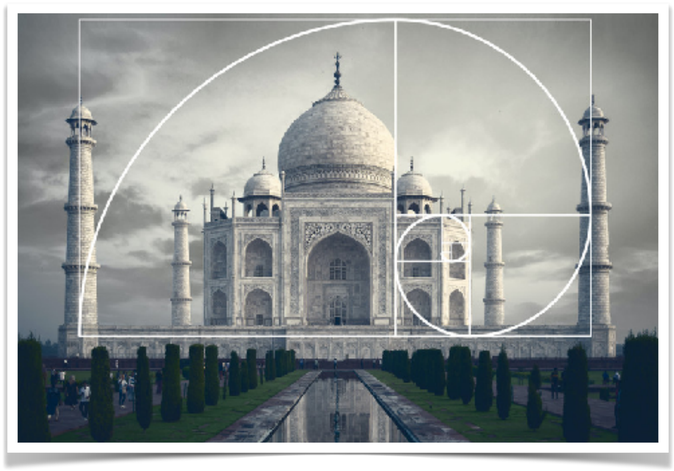

Architecture

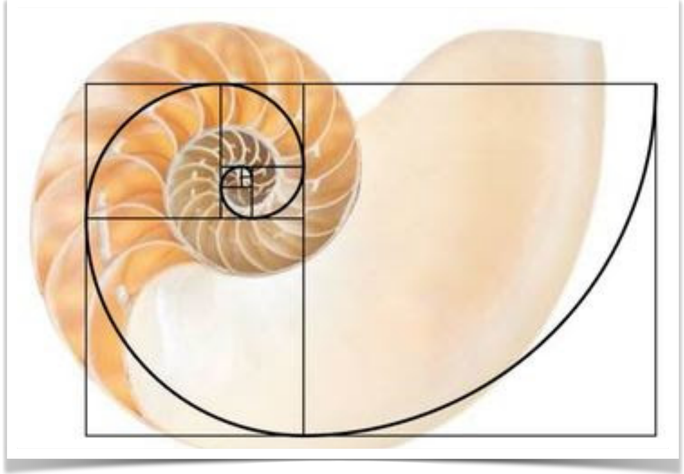

Nature

And now also in computer science :)

Weird Properties

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

Expand Array by Golden Ratio (G = ϕ = 1.61…)

Expand Array by Golden Ratio (G = ϕ = 1.61…)

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 = ϕ = 1.61…)

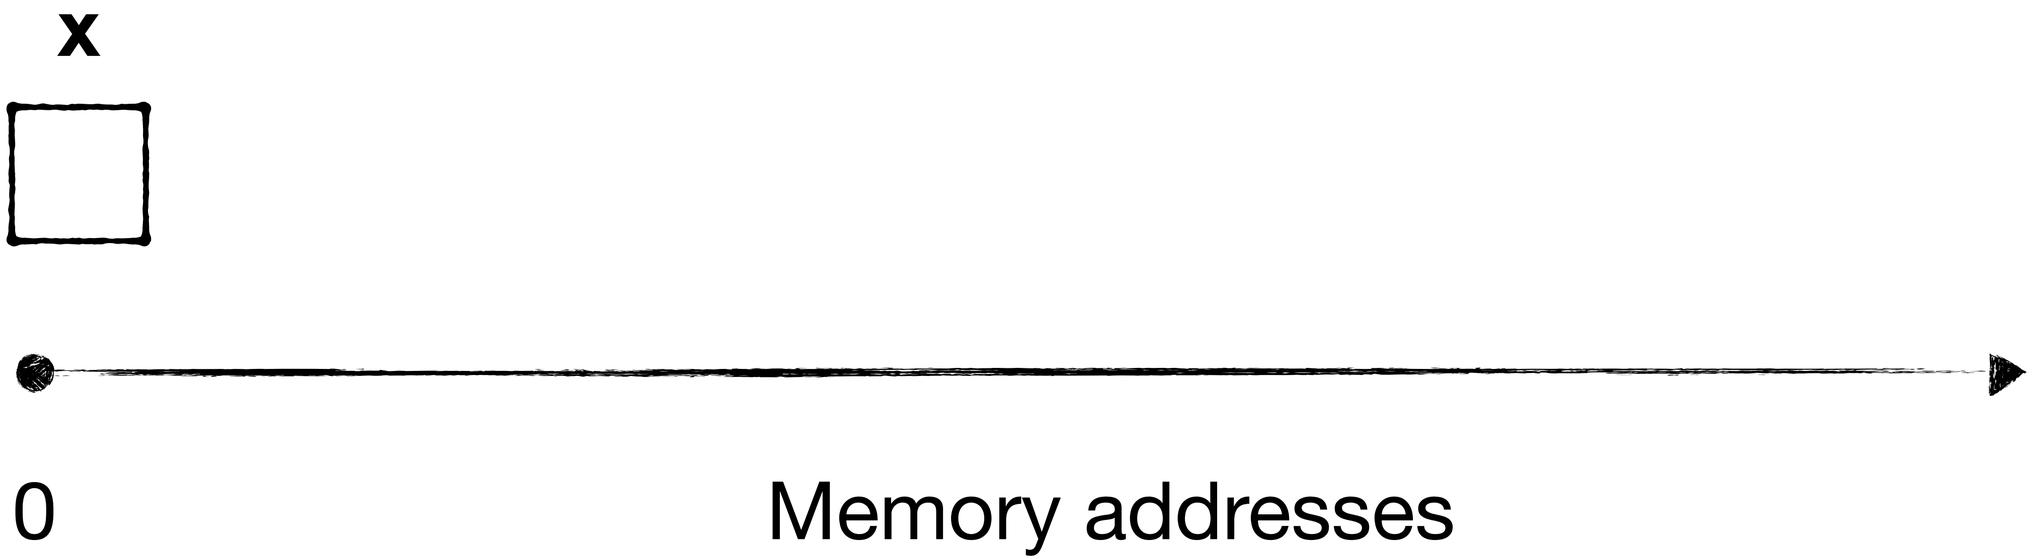

Expand Array by Golden Ratio (G = ϕ = 1.61…)

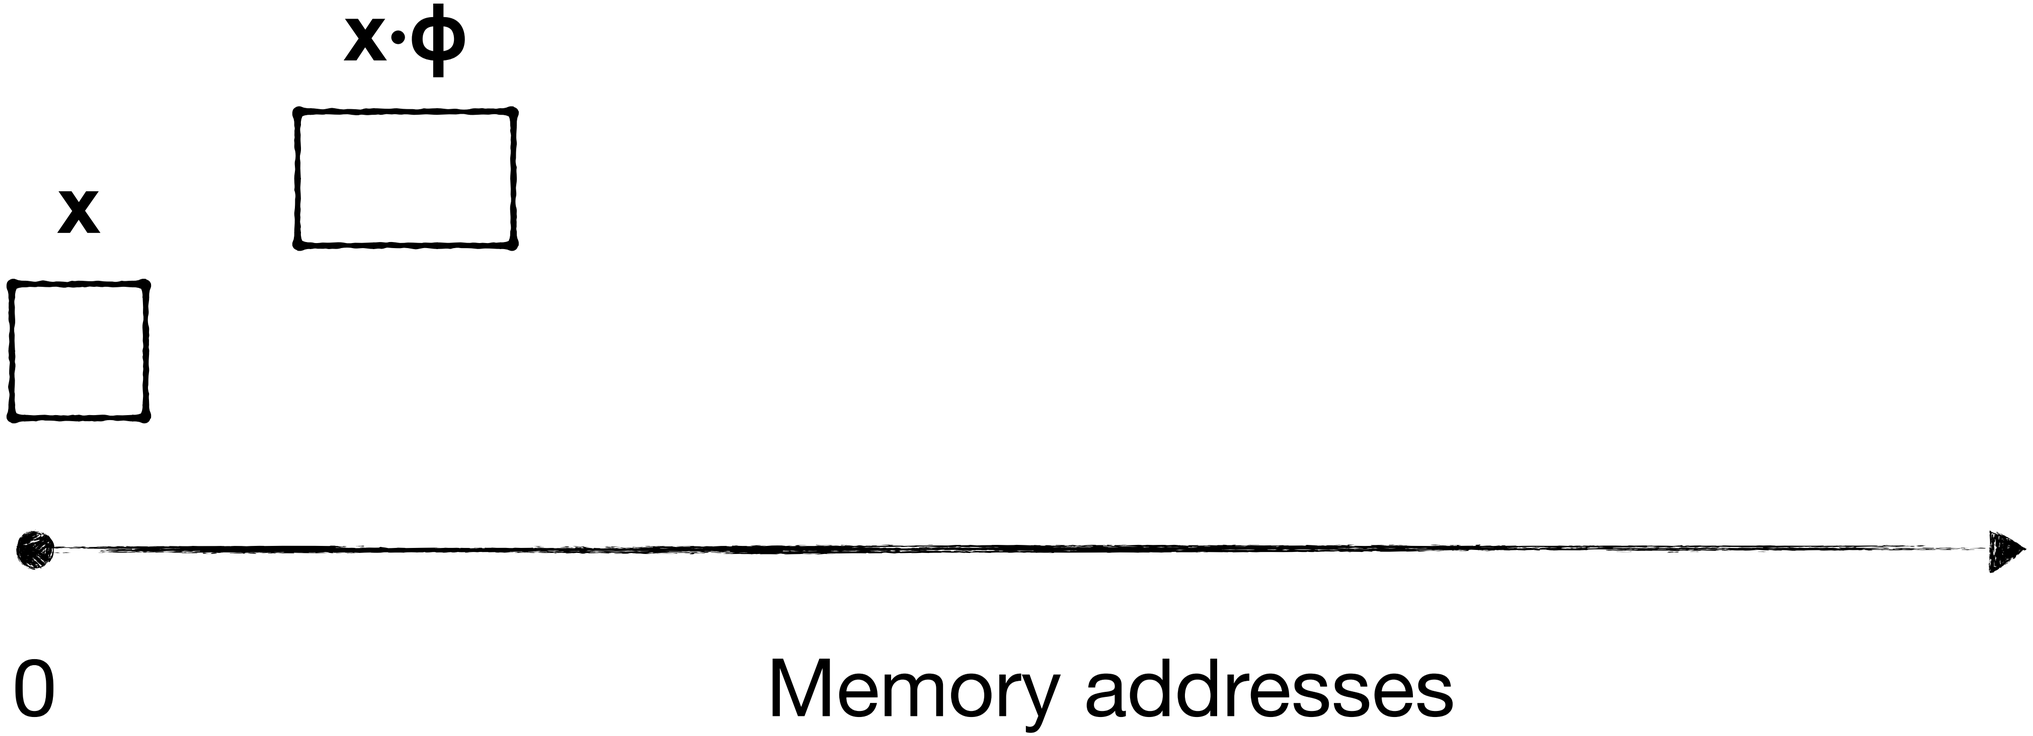

Expand Array by Golden Ratio (G = ϕ = 1.61…)

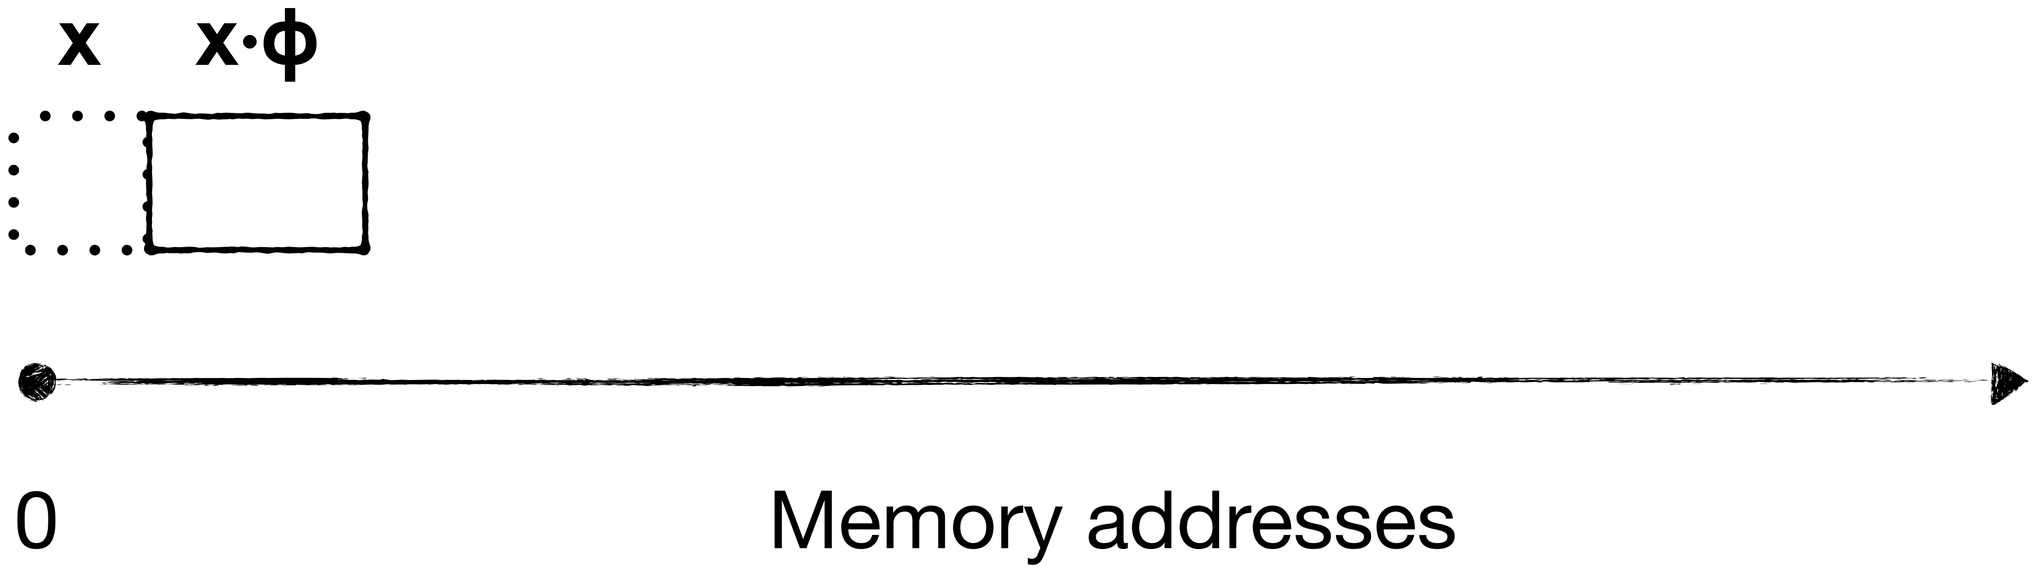

Expand Array by Golden Ratio (G = ϕ = 1.61…)

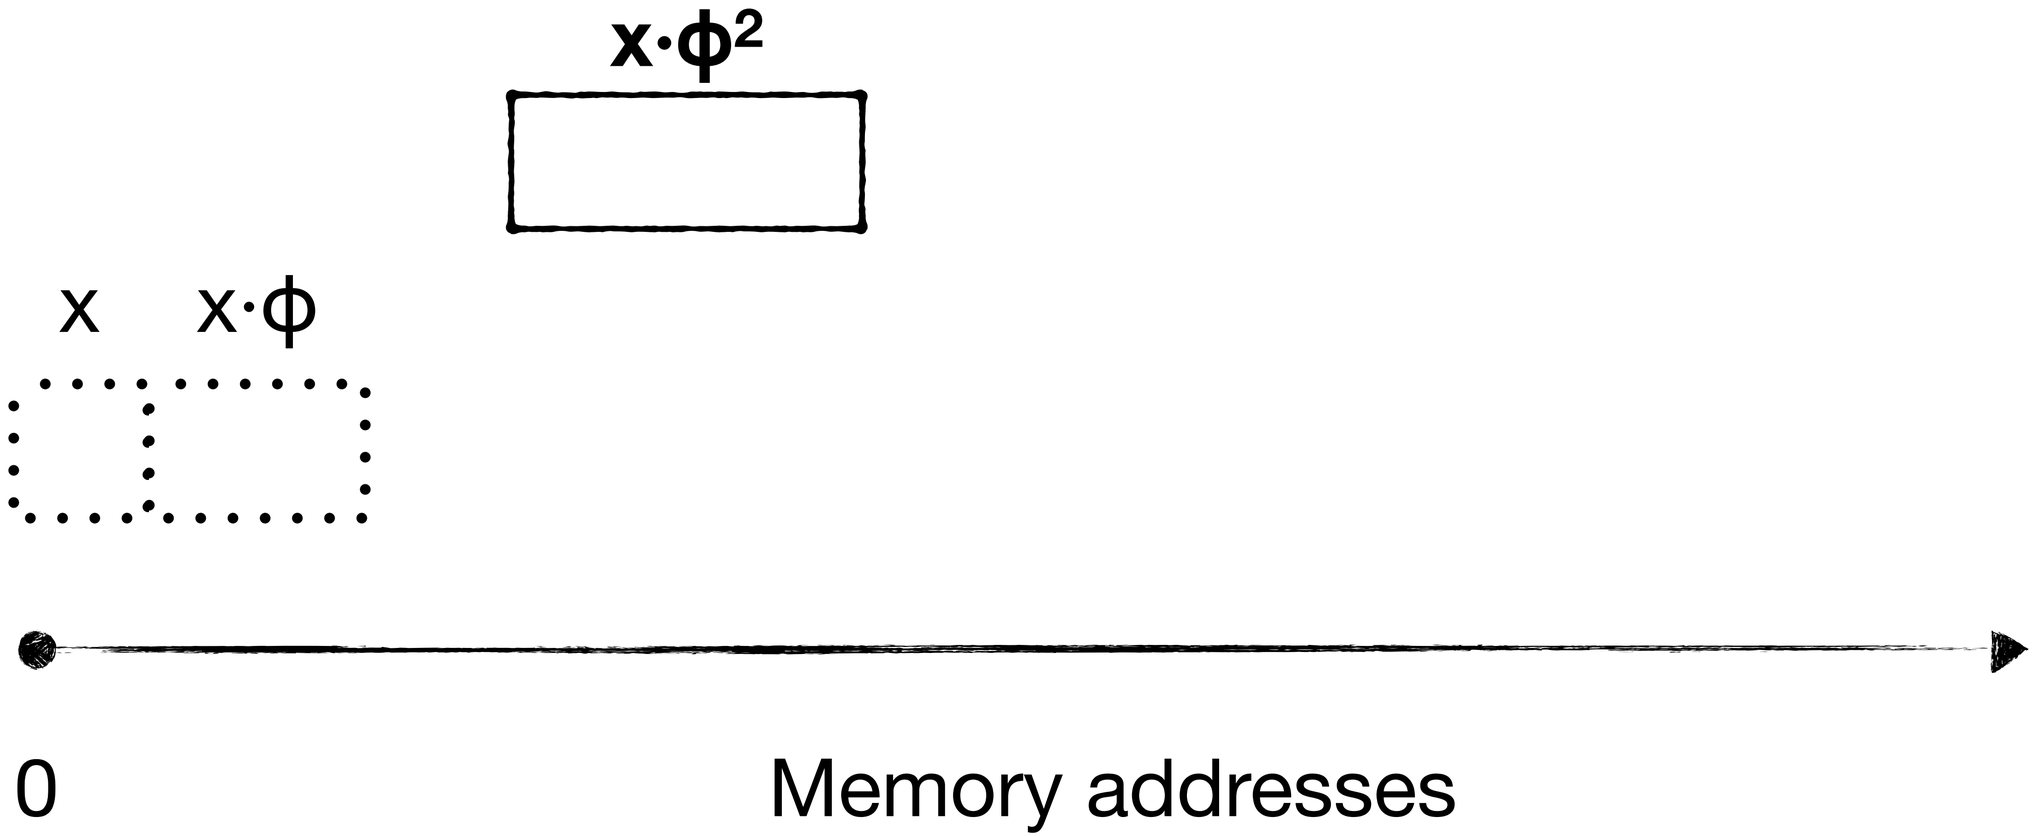

Expand Array by Golden Ratio (G = ϕ = 1.61…)

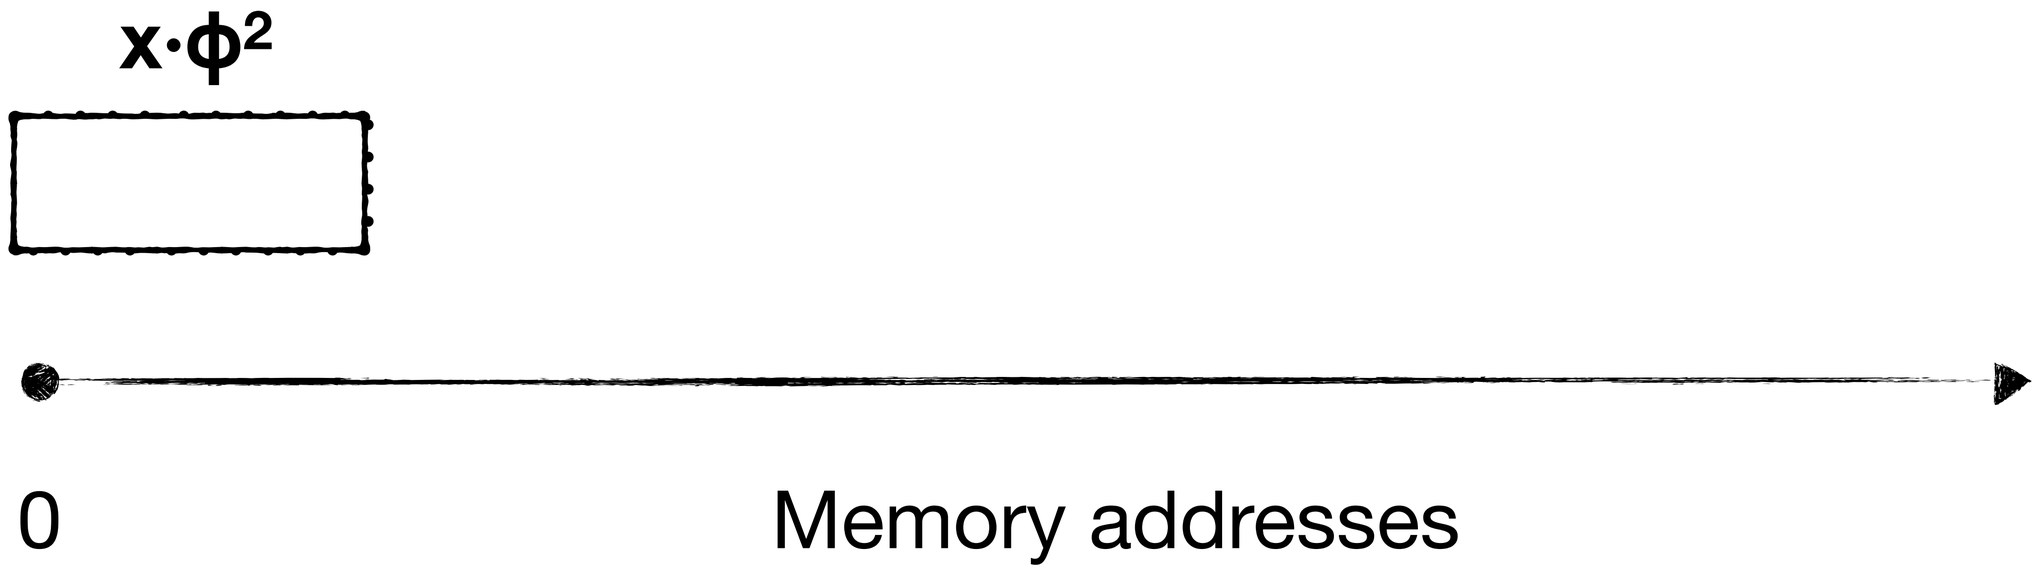

Expand Array by Golden Ratio (G = ϕ = 1.61…)

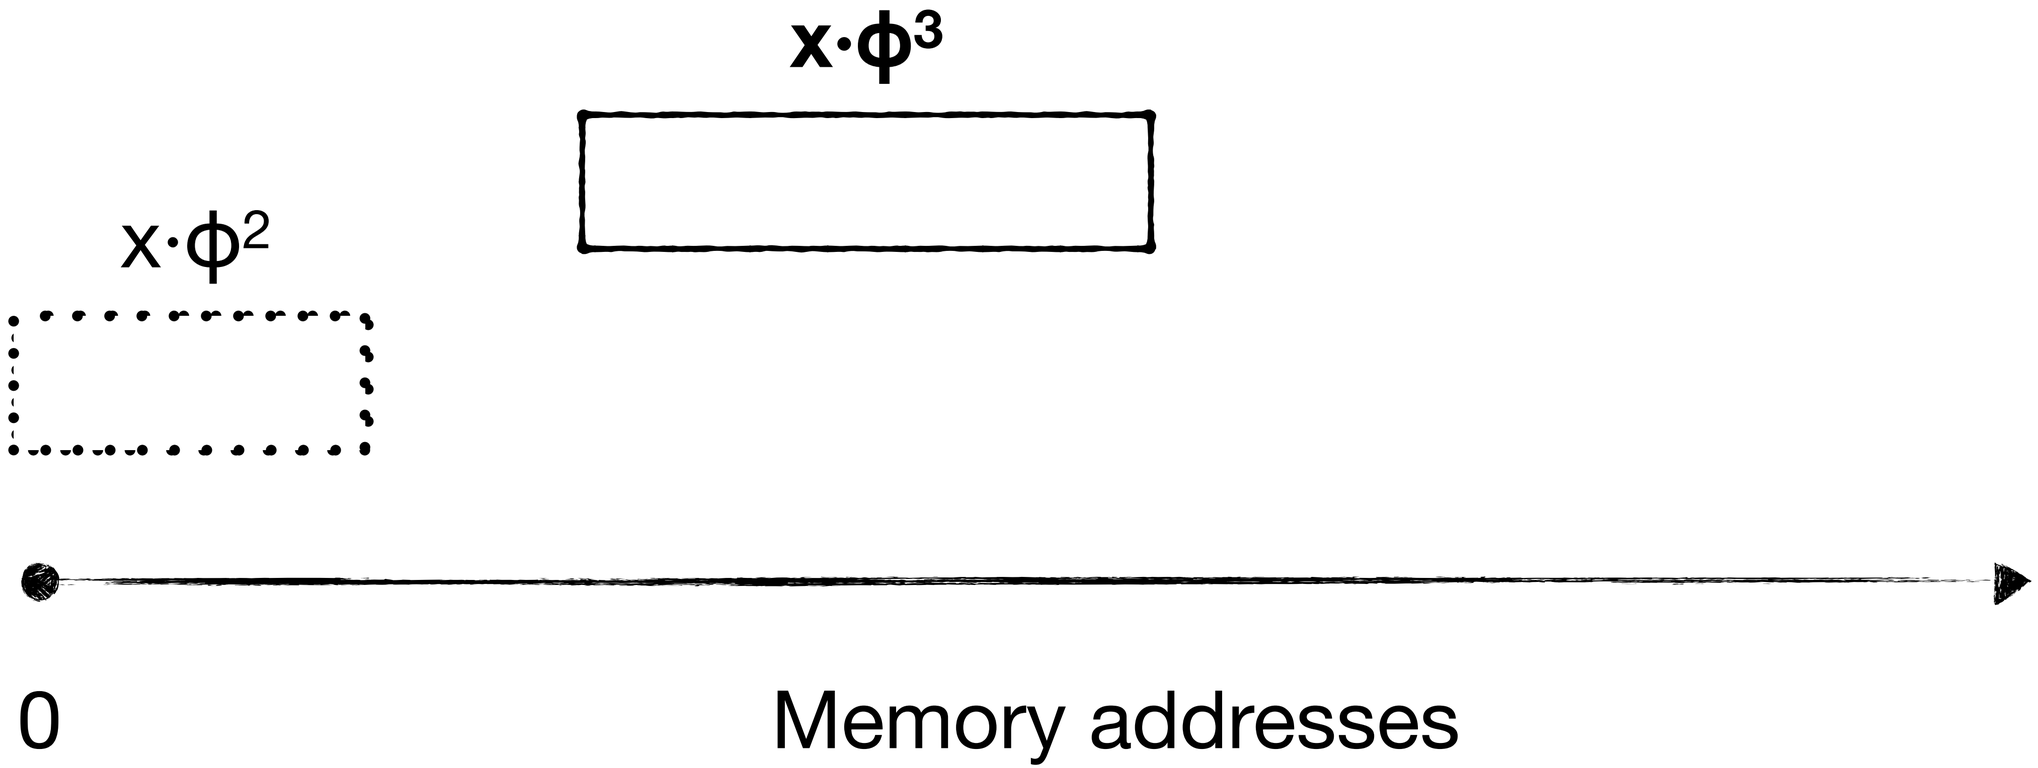

Expand Array by Golden Ratio (G = ϕ = 1.61…)

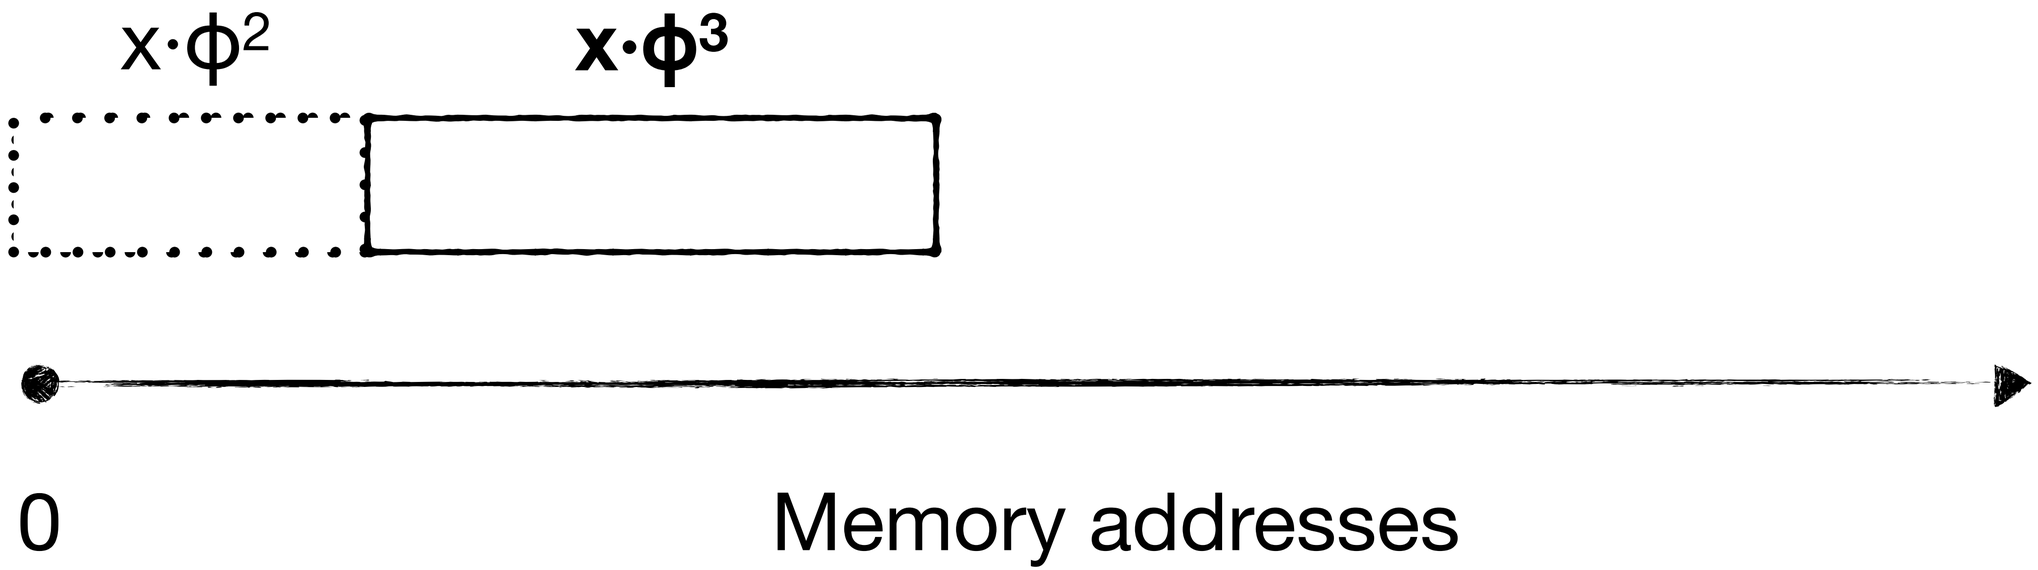

Expand Array by Golden Ratio (G = ϕ = 1.61…)

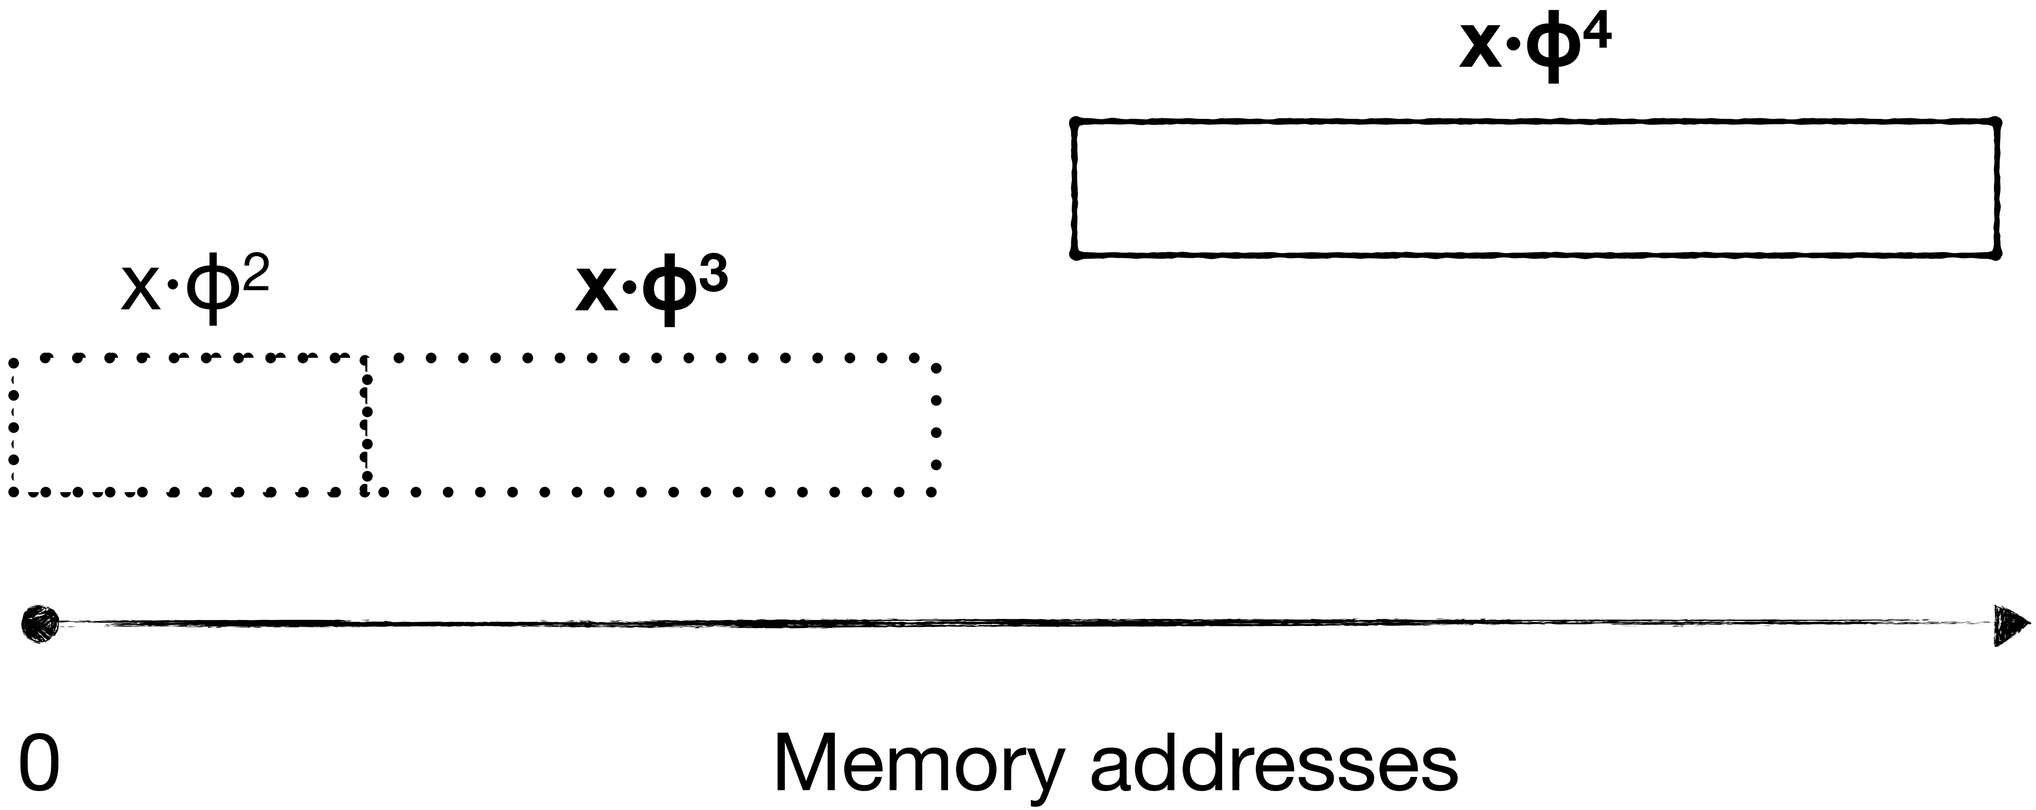

Expand Array by Golden Ratio (G = ϕ = 1.61…)

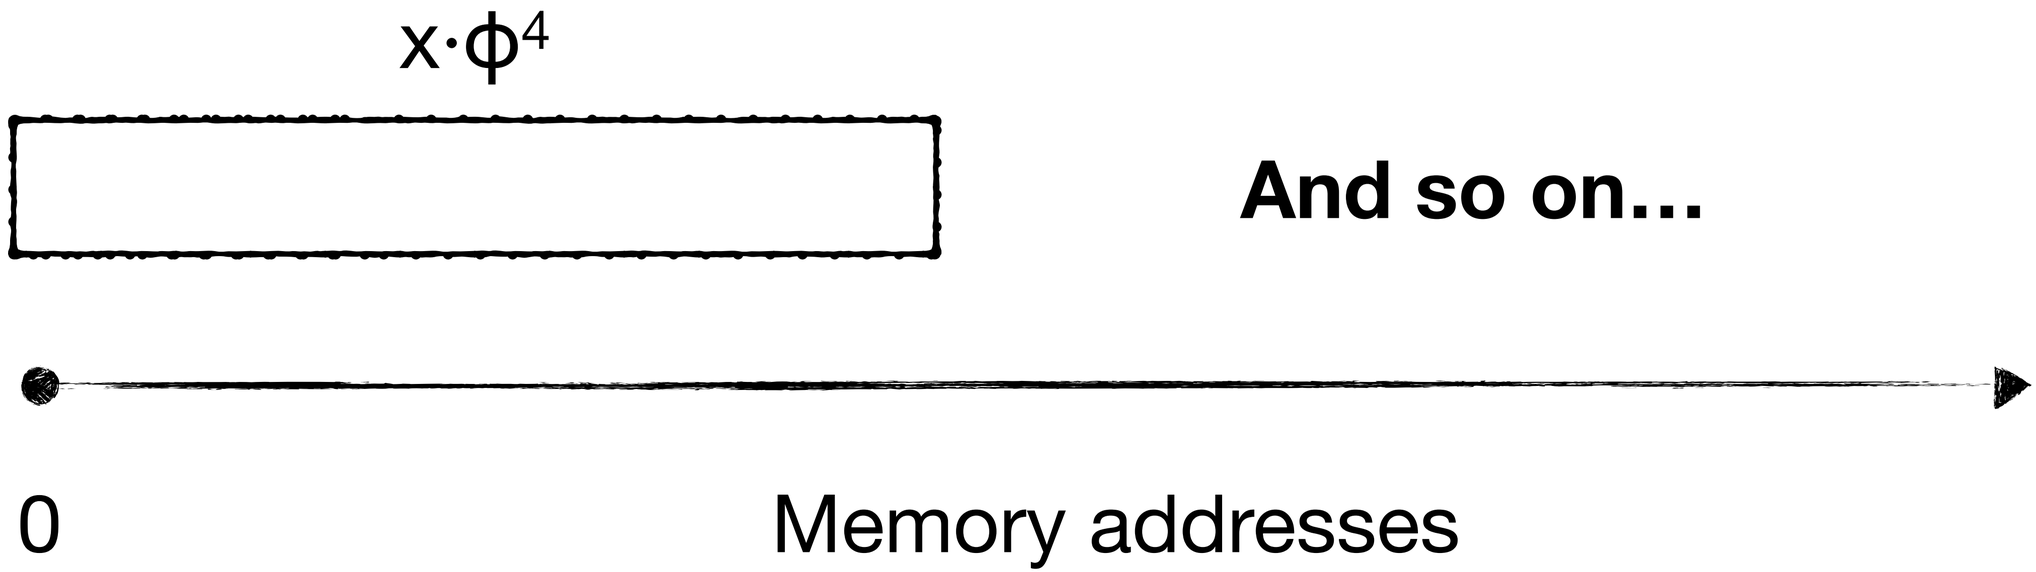

Write-Amplification

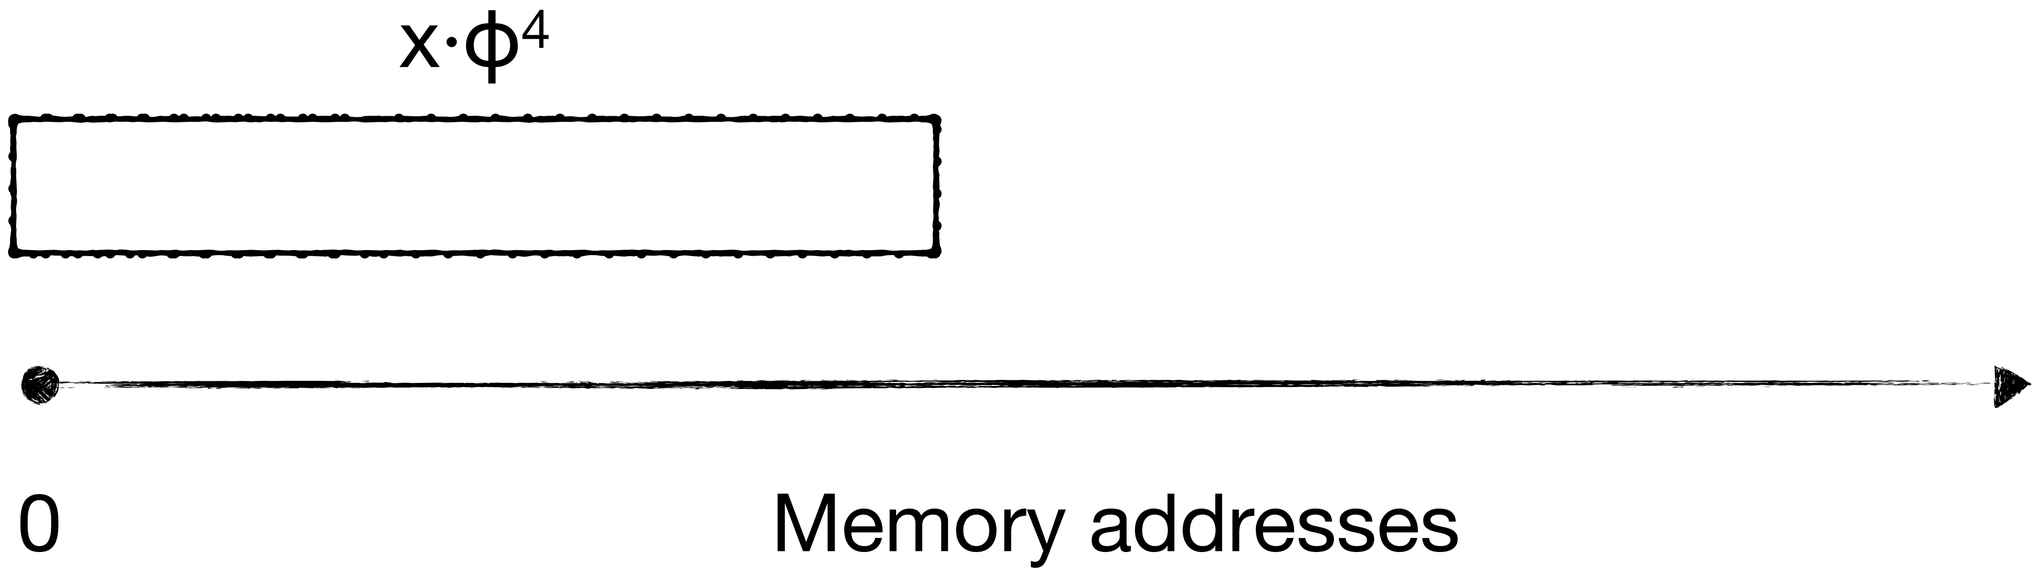

$$\textbf{Write-Amplification} = \frac{\textbf{G}}{\textbf{G - 1}}$$

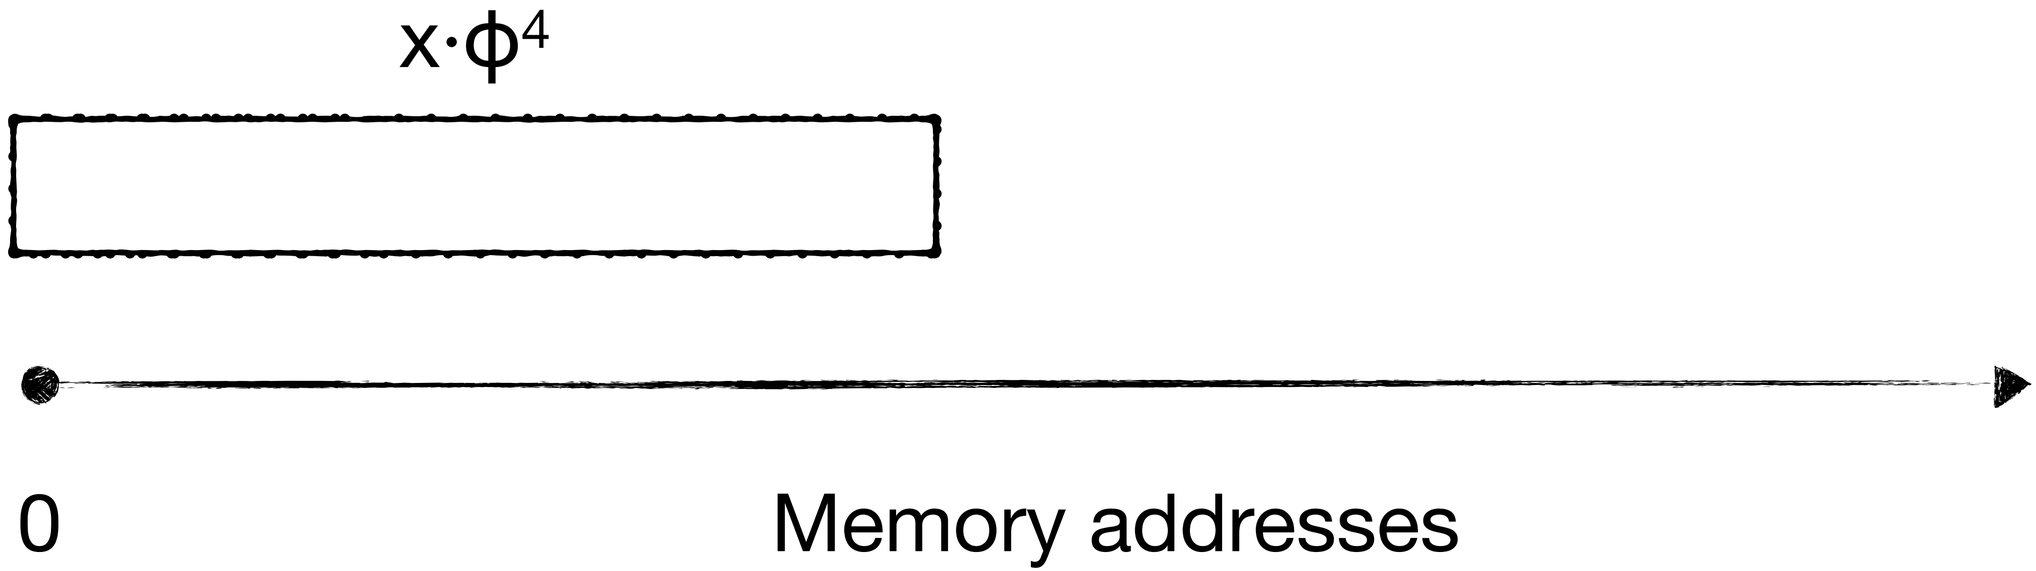

$$\text{Write-Amplification} = \frac{G}{G - 1} = \frac{\phi}{\phi - 1}$$

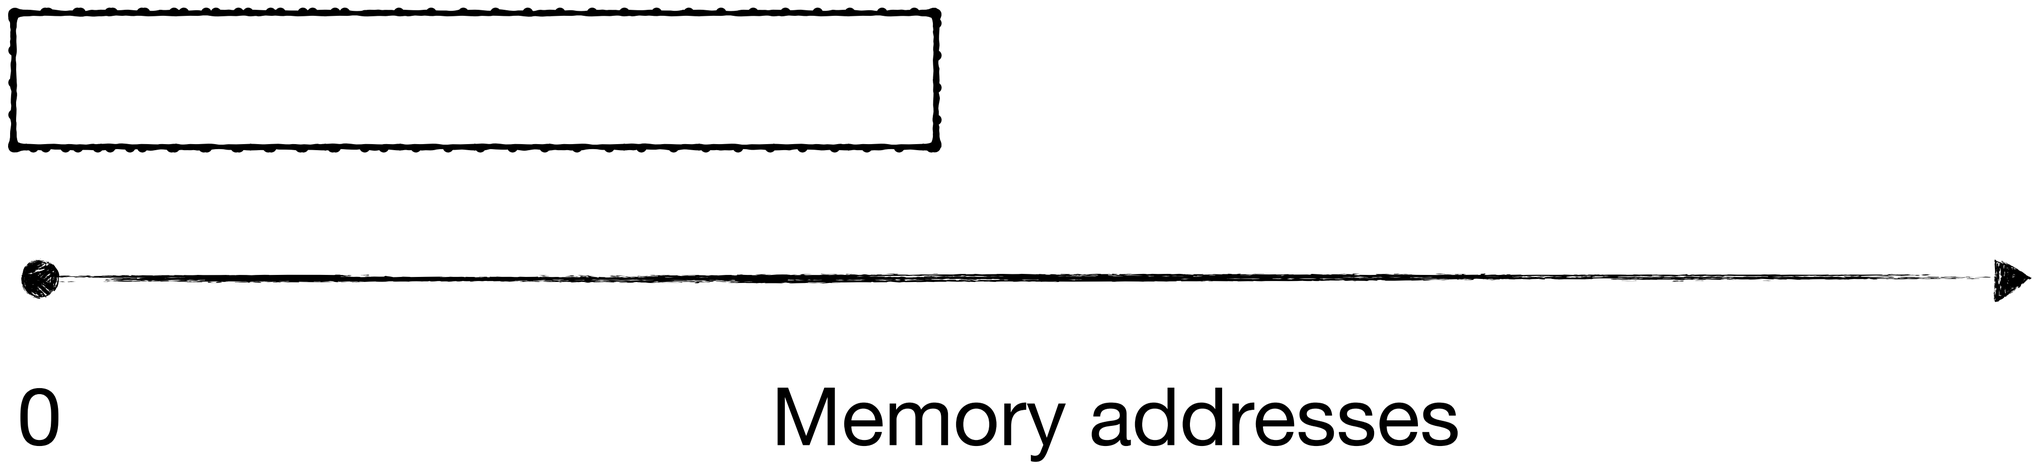

$$\text{Write-Amplification} = \frac{G}{G - 1} = \frac{\phi}{\phi - 1} = \phi + 1$$

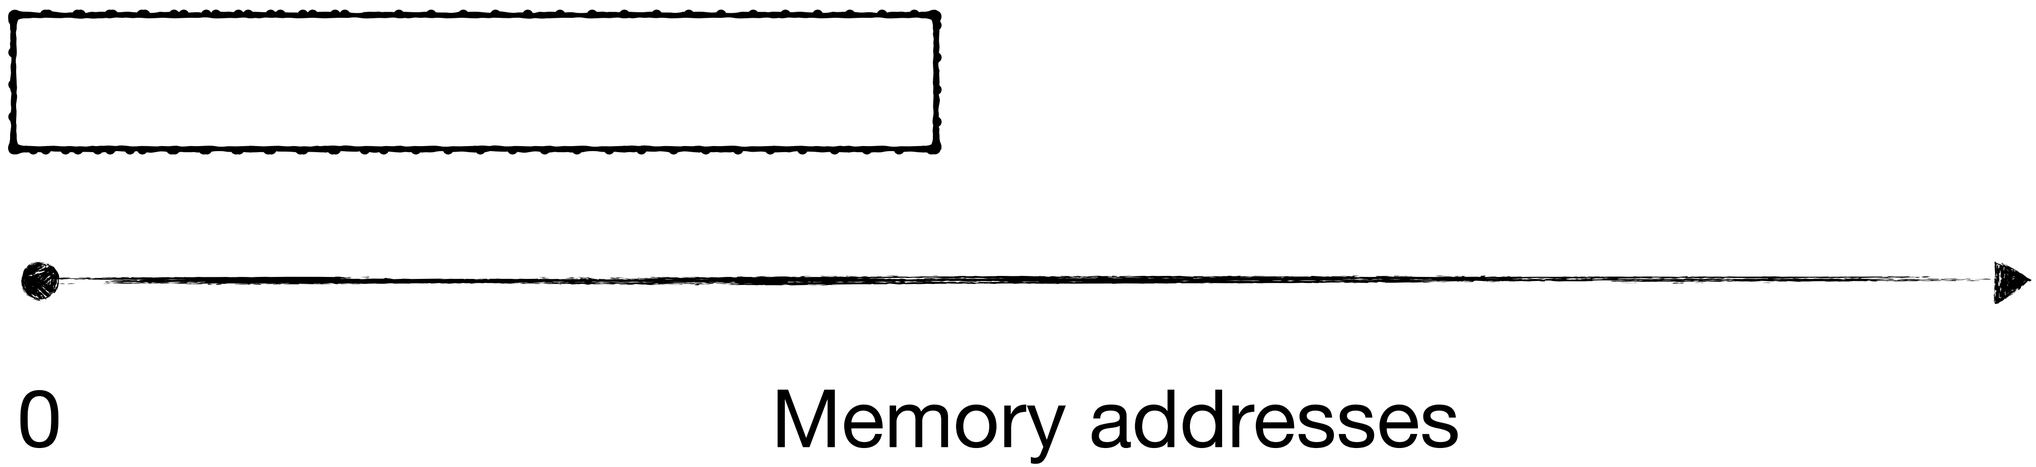

Space-Amplification?

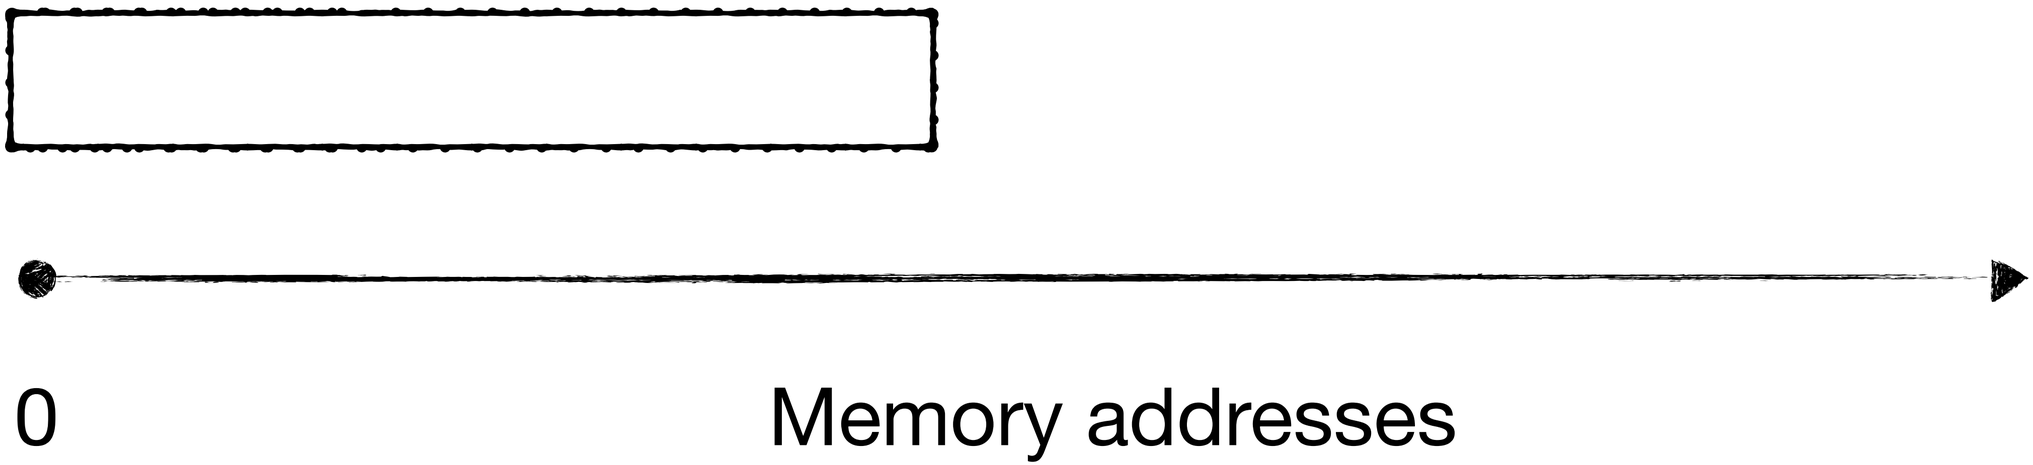

Space-Amplification?

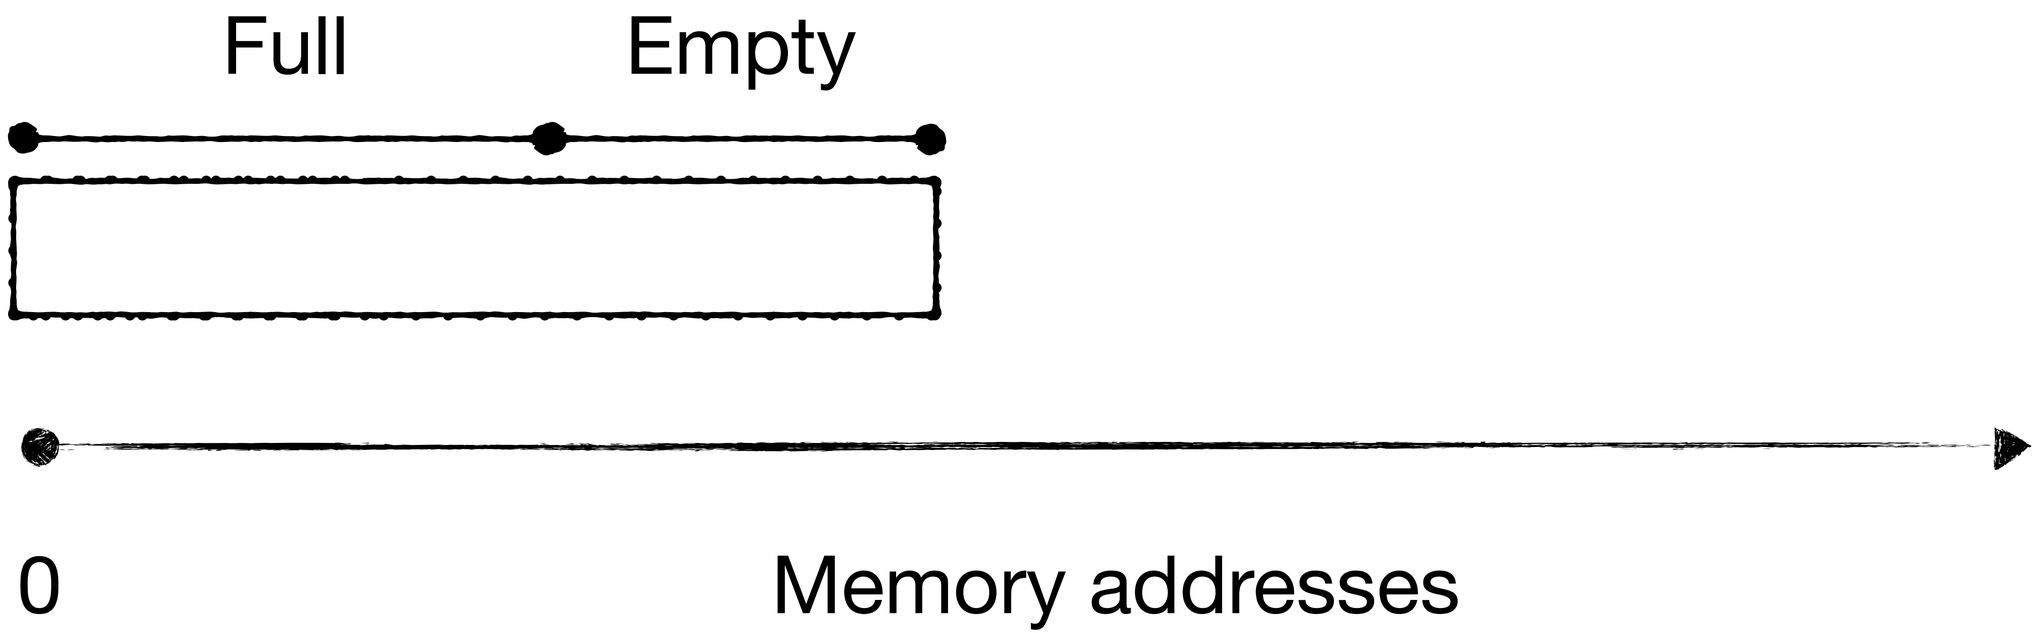

$$\text{Space-Amplification} = \frac{\textbf{Full + Empty}}{\textbf{Full}} = \textbf{G} = \boldsymbol{\phi}$$

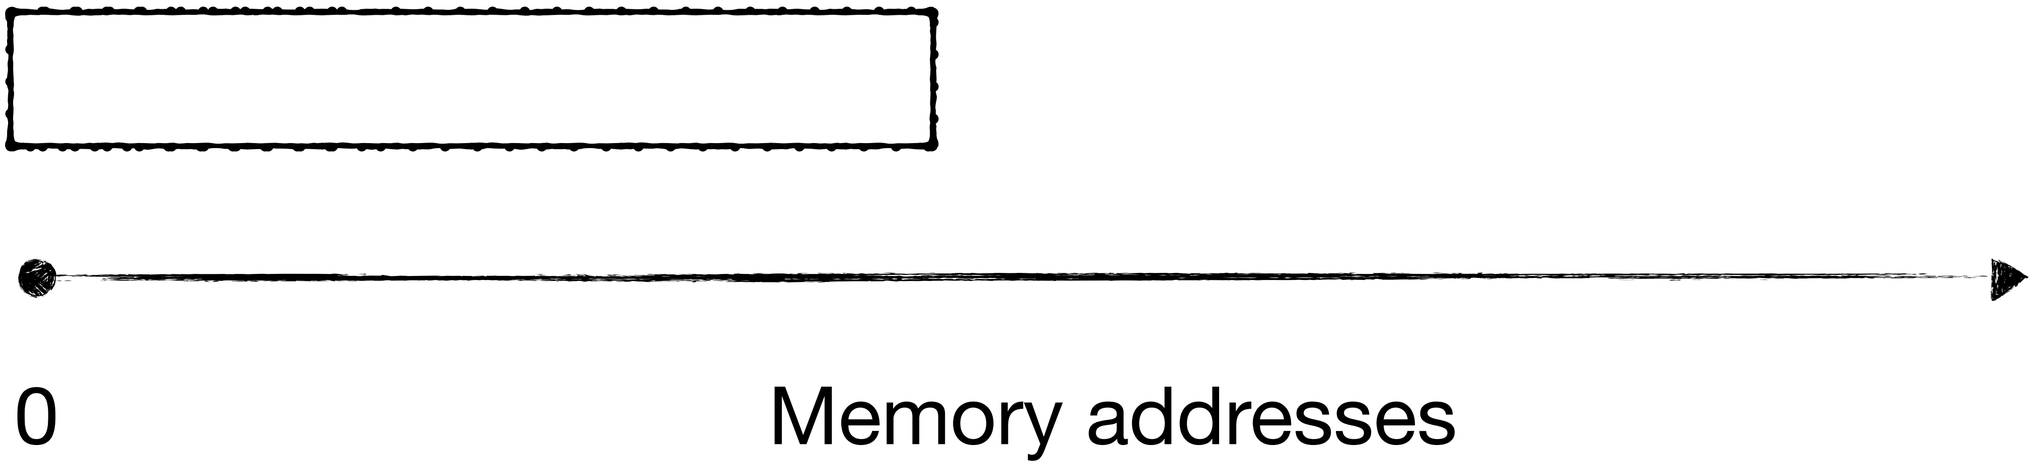

Max Space-Amp?

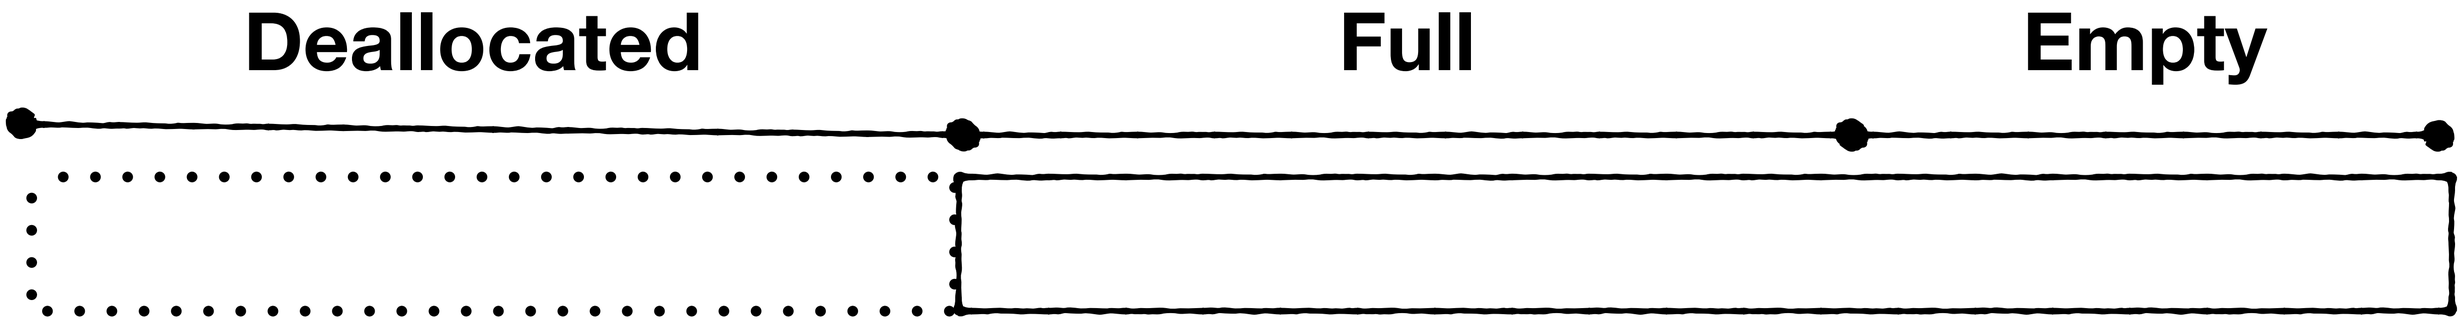

For G < ϕ

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

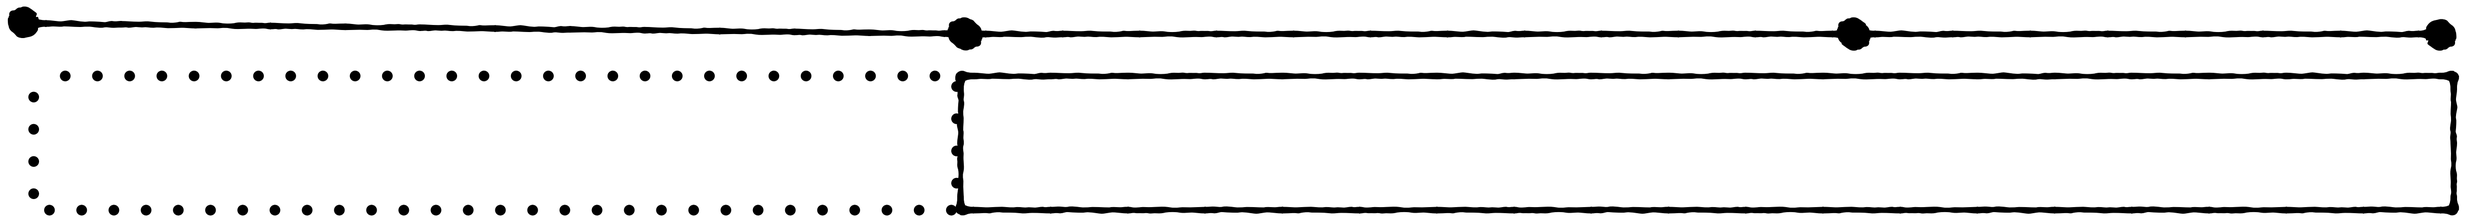

Max Space-Amp = **1 + ϕ** = **2.61**

	Write-amp	Space-amp
$G > \phi$	$\frac{G}{G - 1}$	$\frac{G}{G - 1} + G$
$G < \phi$	$\frac{G}{G - 1}$	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

And now to new stuff

Reusing Deallocated Space

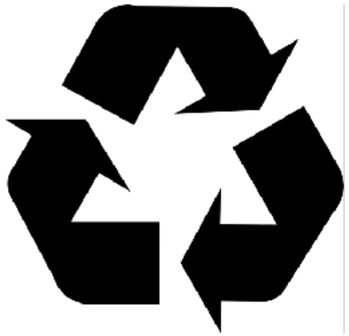

Alleviating trade-off via indirection

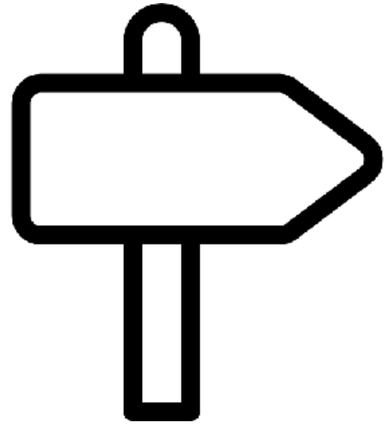

Can we completely overcome this trade-off?

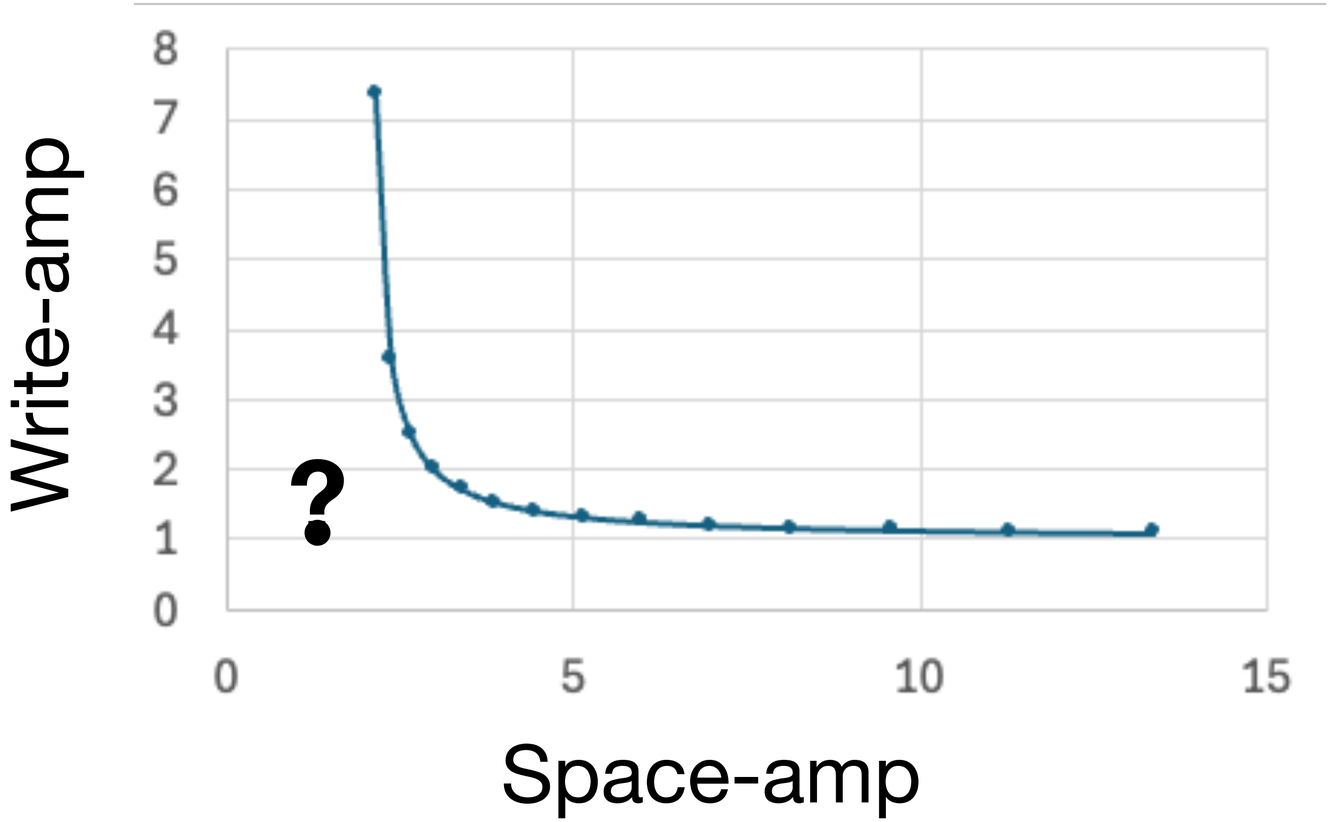

Suppose we could expand without copying everything:

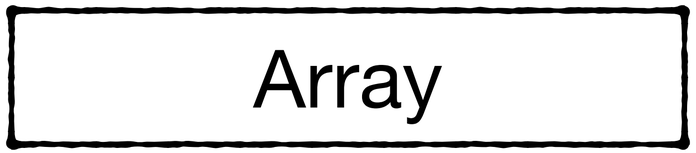

Suppose we could expand without copying everything:

Suppose we could expand without copying everything:

Promise: **write-amp of ???**

Suppose we could expand without copying everything:

Promise: **write-amp of 1**

Suppose we could expand without copying everything:

Promise: **write-amp of 1**

space-amp? We shall see :)

Add a layer of indirection

Directory

Data blocks

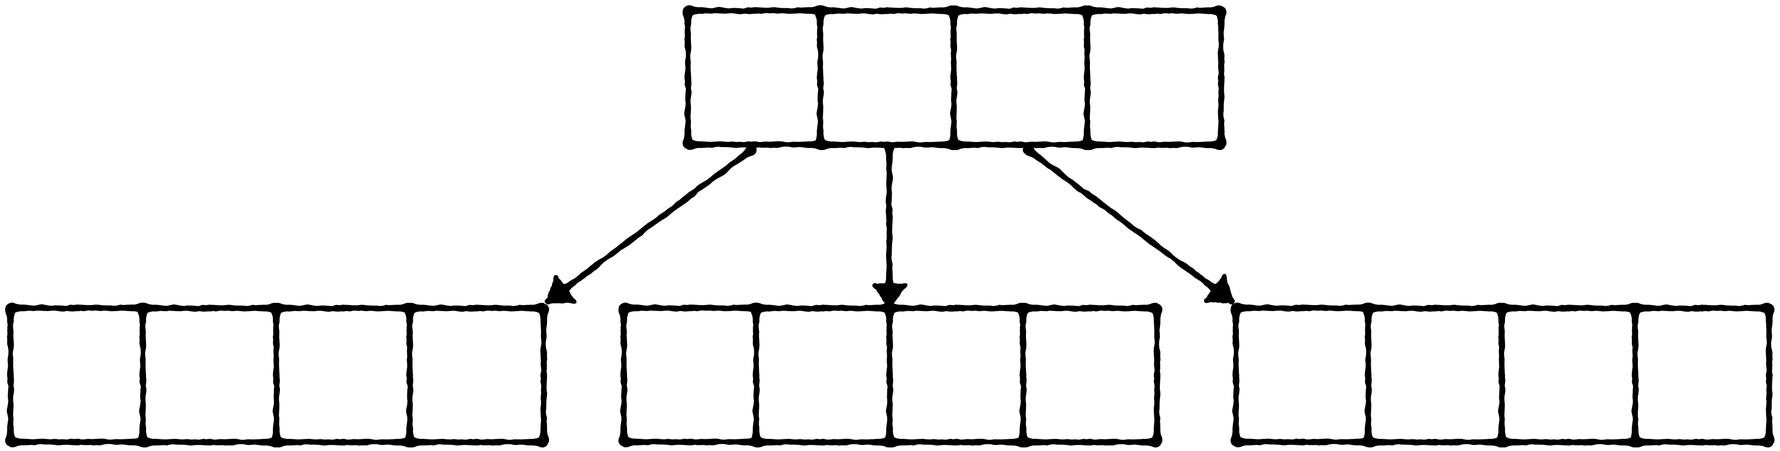

get(i)

Directory

Data blocks

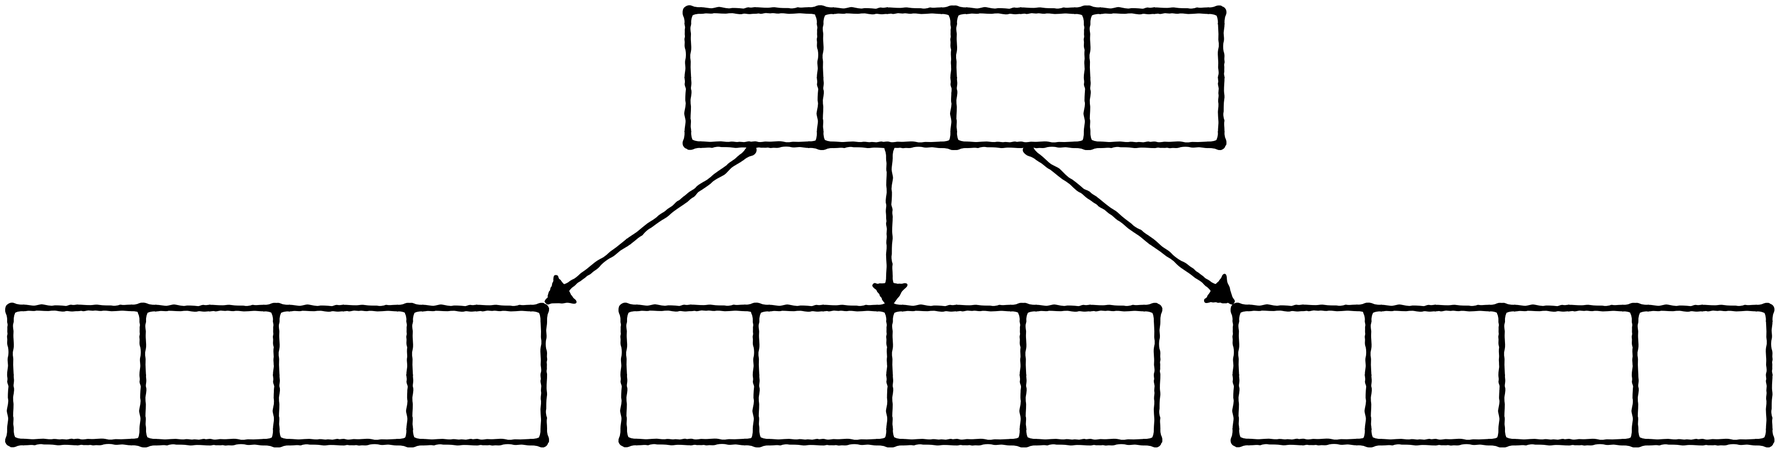

get(i) **Data block = ⌊i / data block size⌋**

Directory

Data blocks

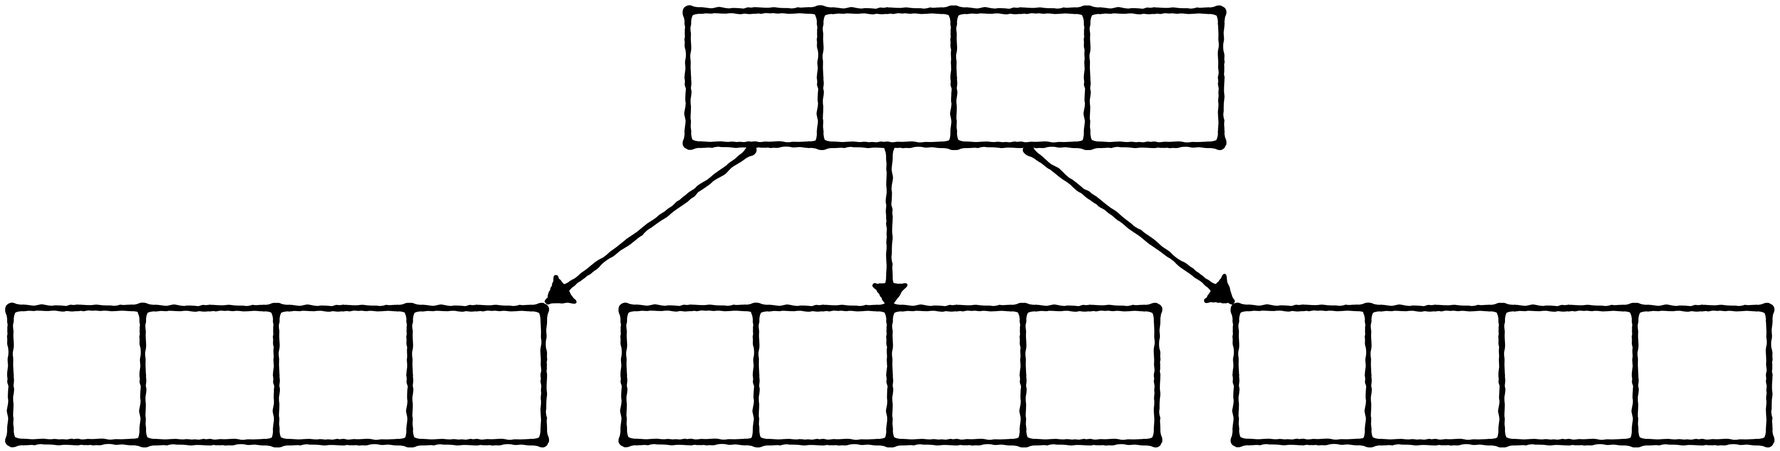

get(i)

Data block = ⌊i / data block size⌋

offset within = i % data block size

Directory

Data blocks

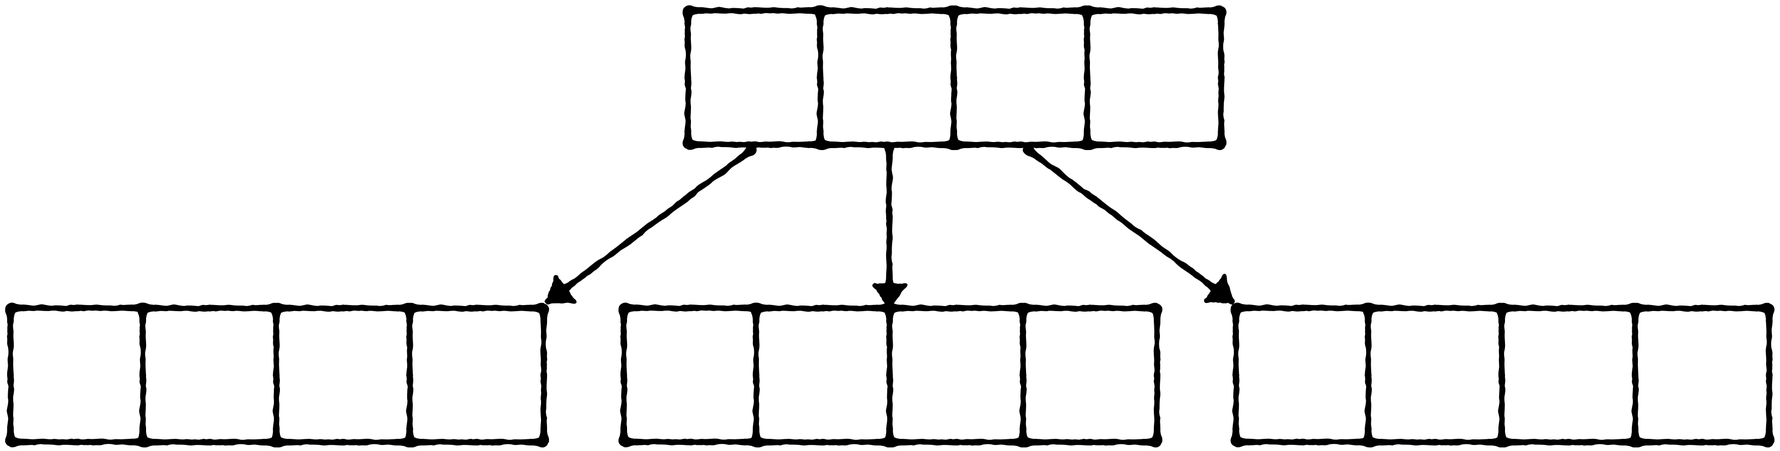

get(5)

Data block = ⌊5 / 4⌋ = **1**

offset within = 5 % 4 = **1**

Directory

Data blocks

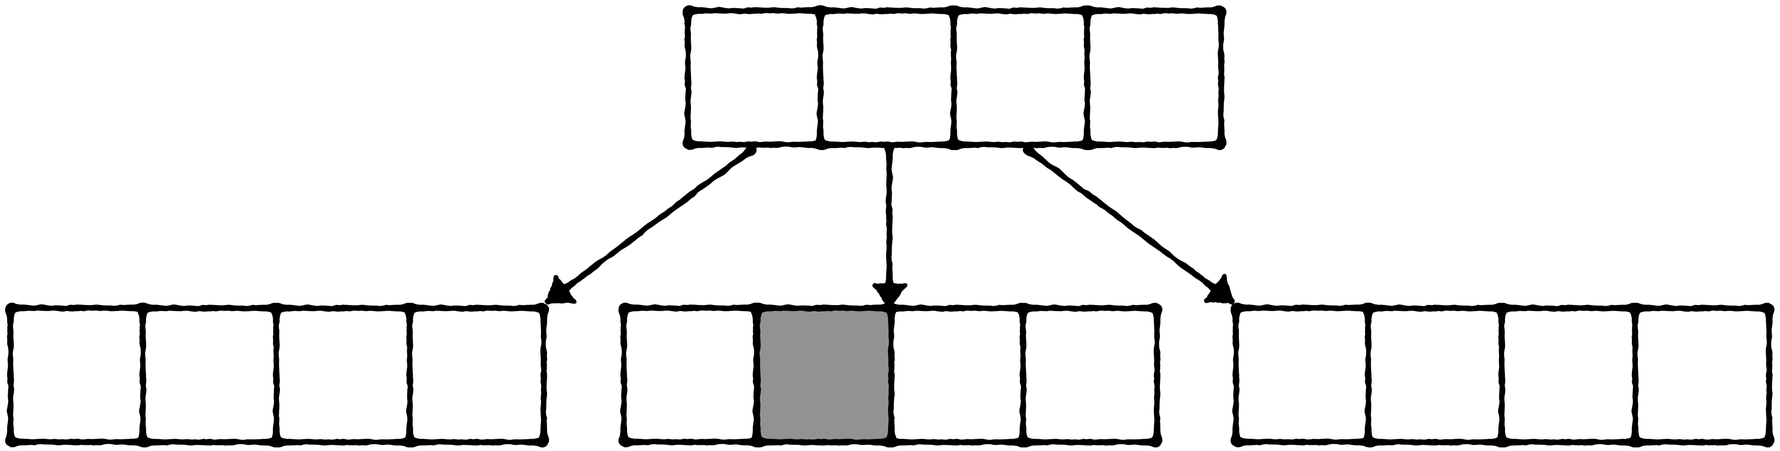

Expand?

Directory

Data blocks

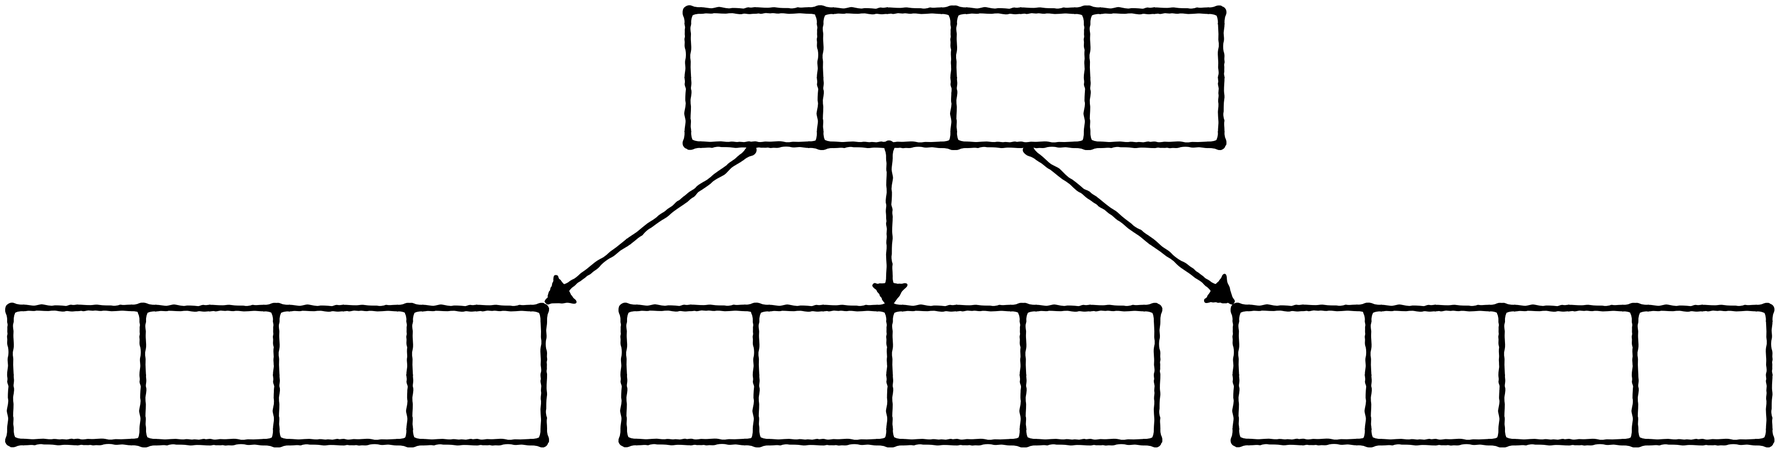

Expand?

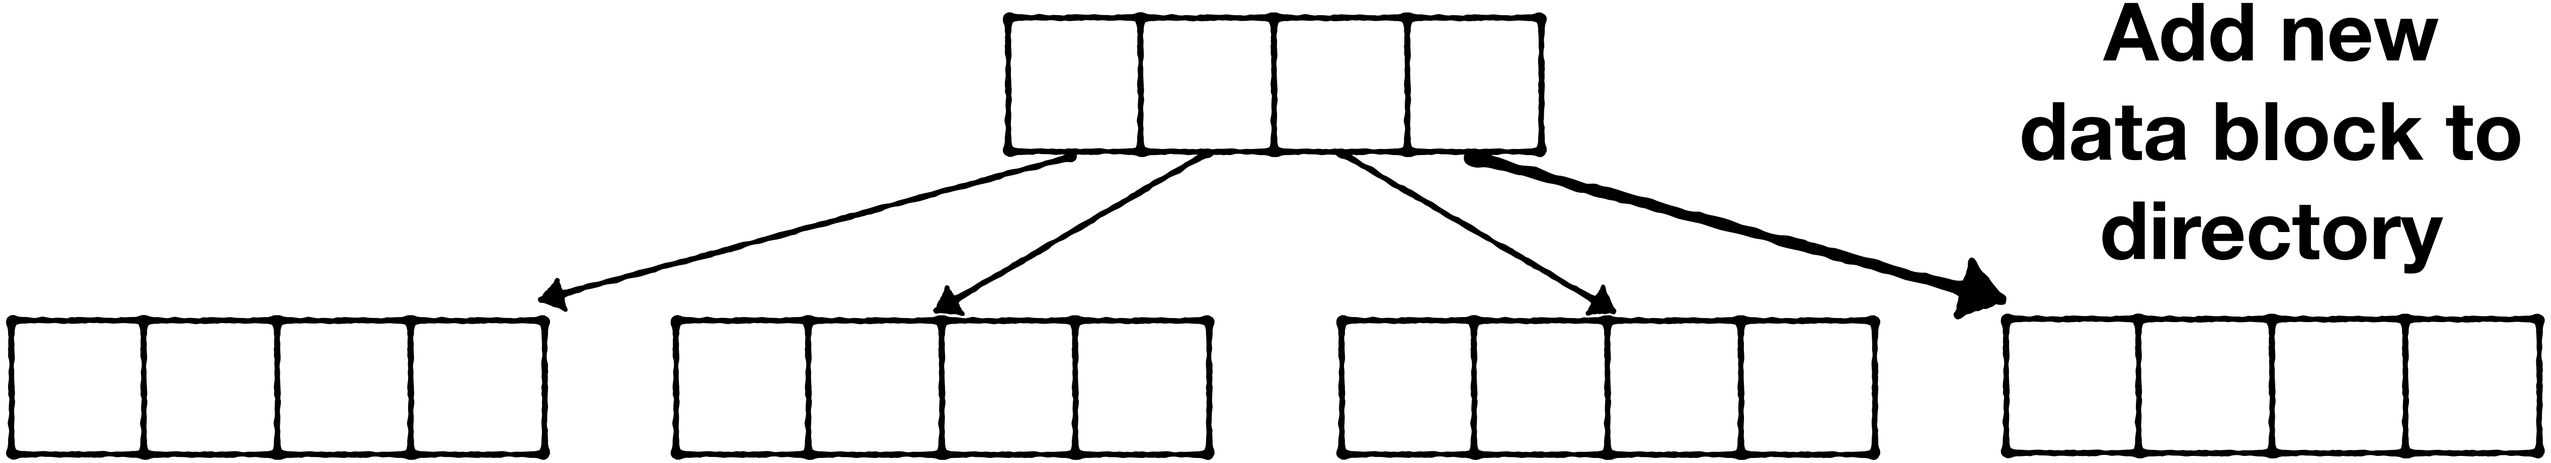

Expand?

Expand directory if we need more space

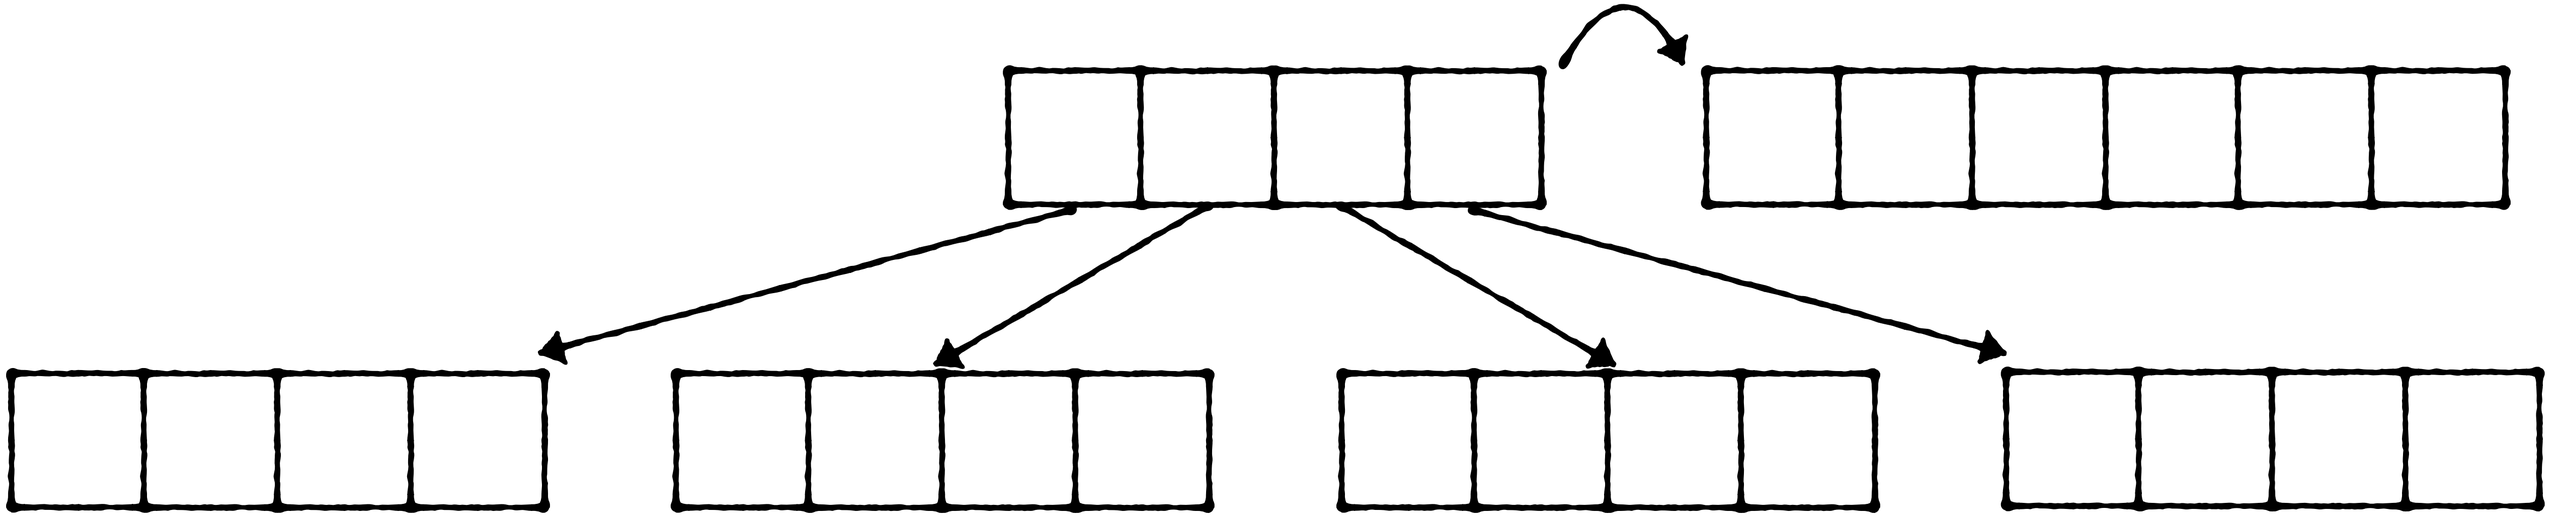

Expand?

Expand directory if we need more space

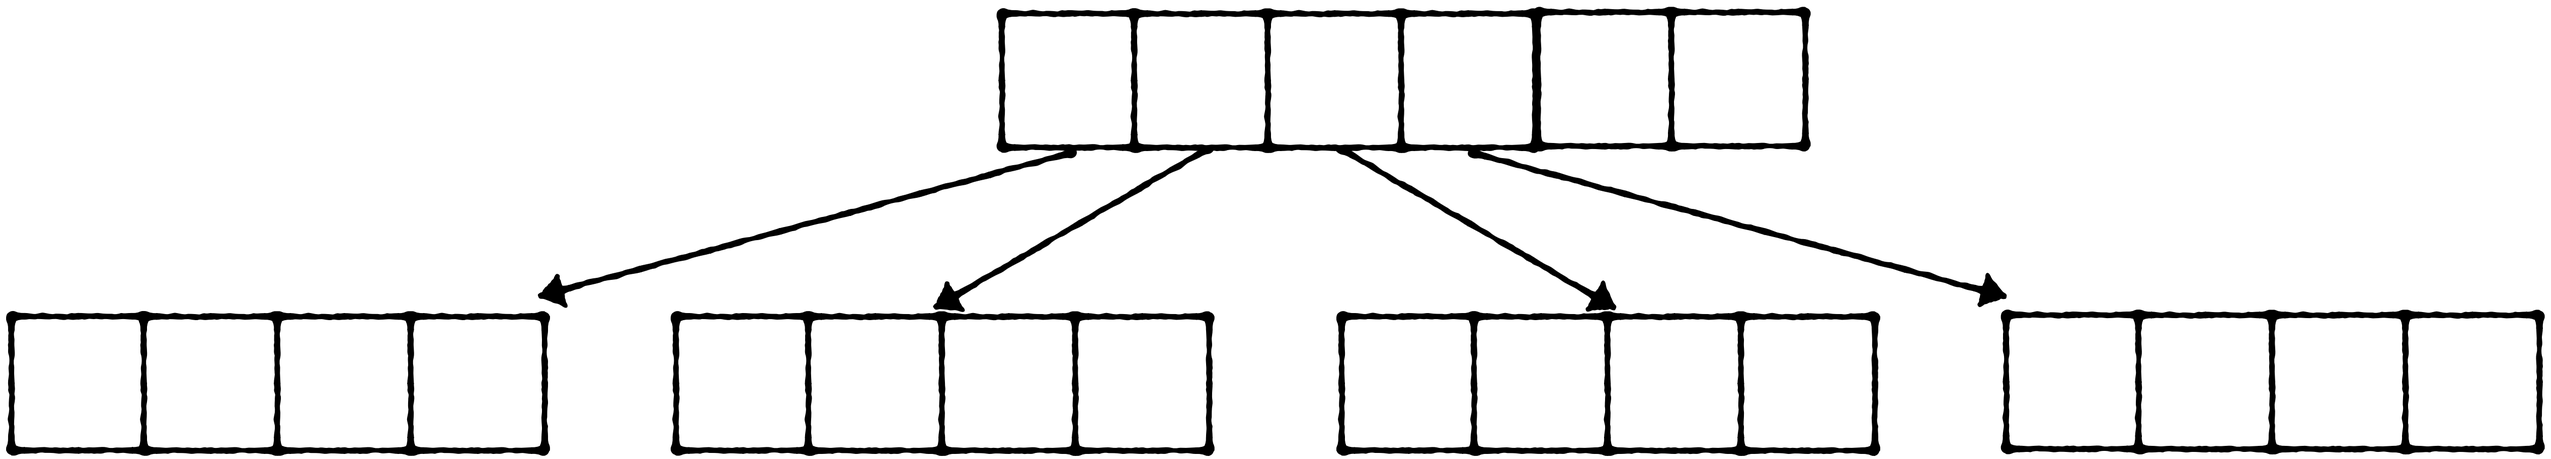

Downside?

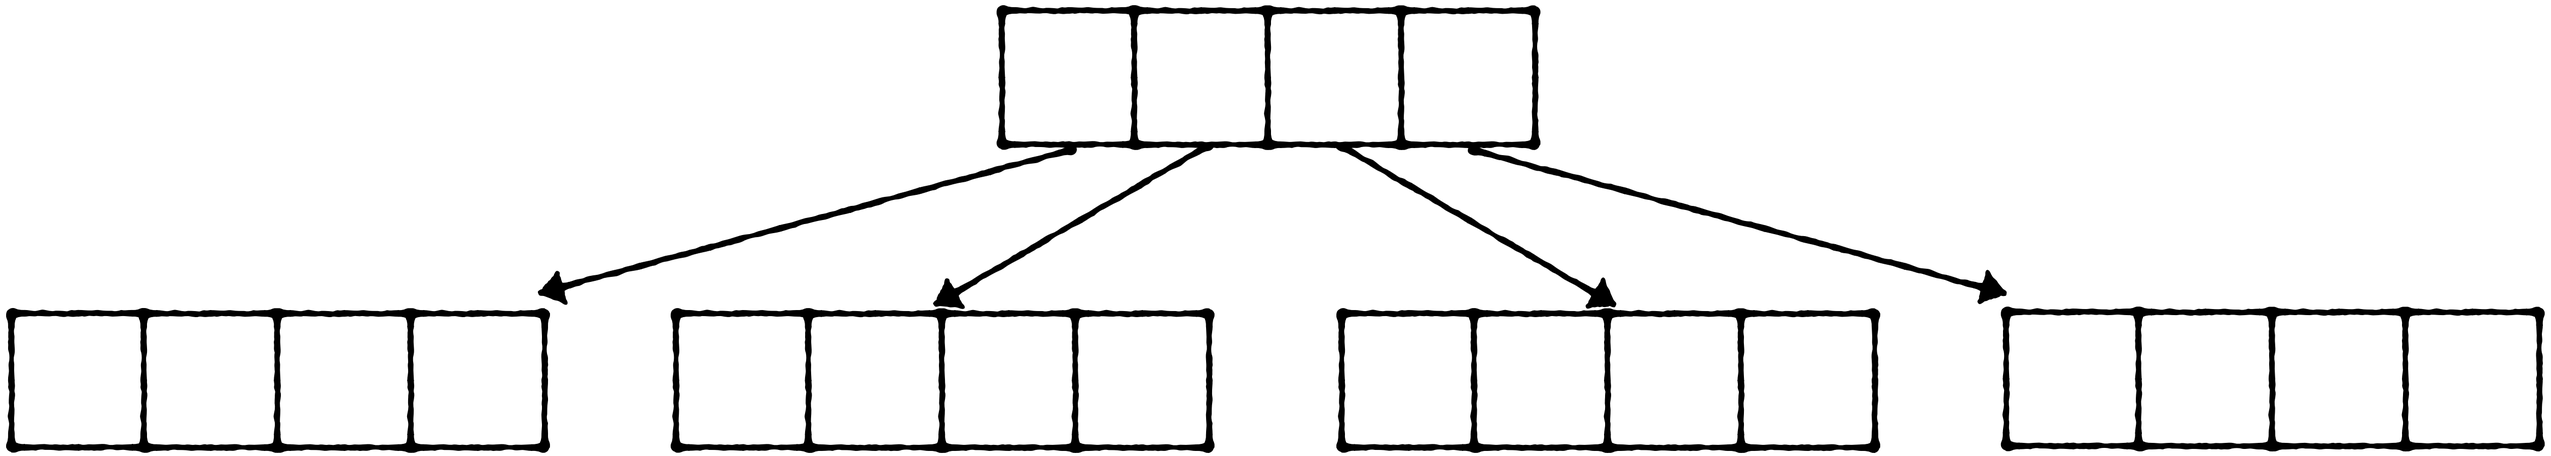

Downside: **2 memory hops per access**

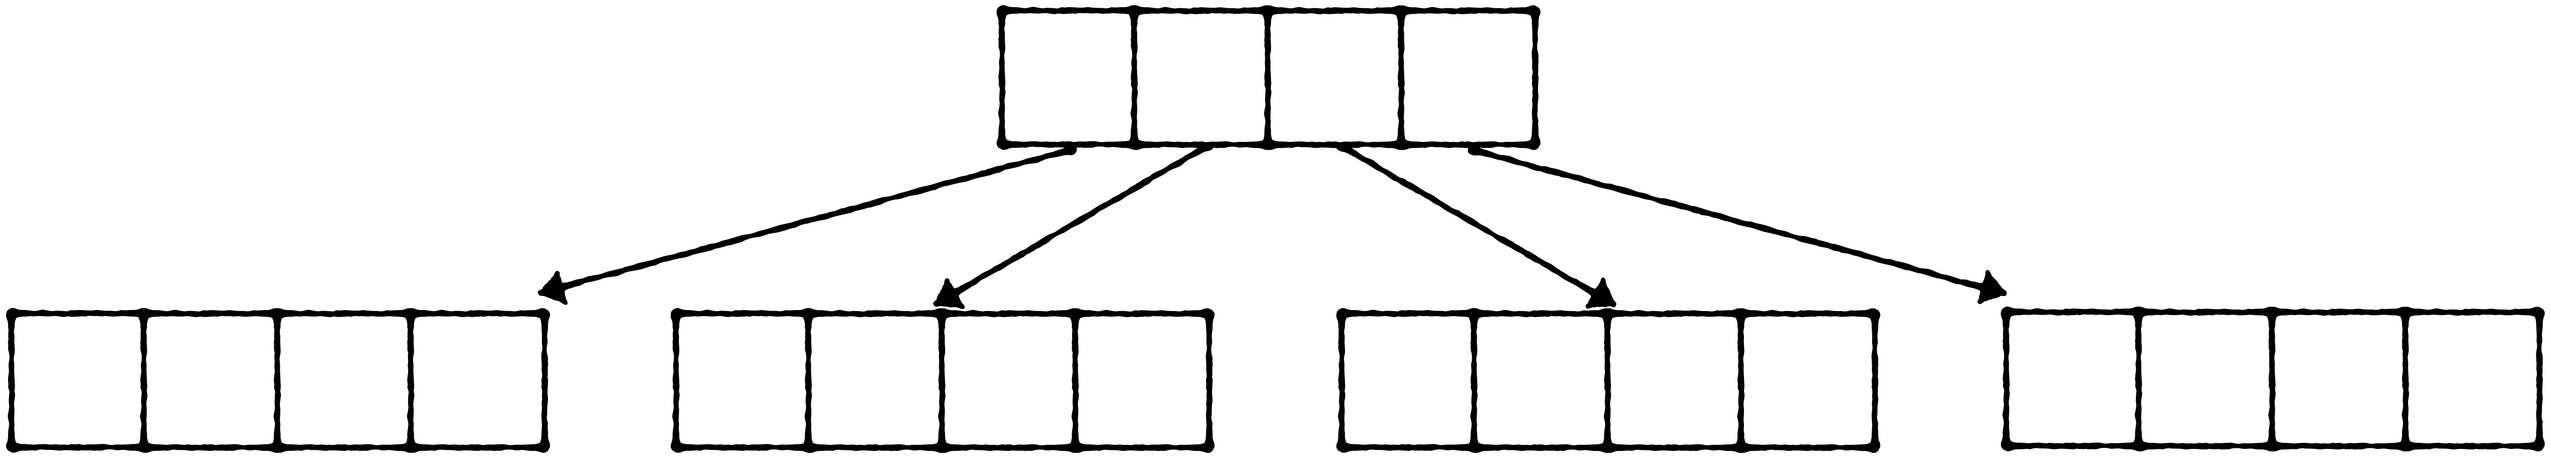

Downside: 2 memory hops per access

Mitigation?

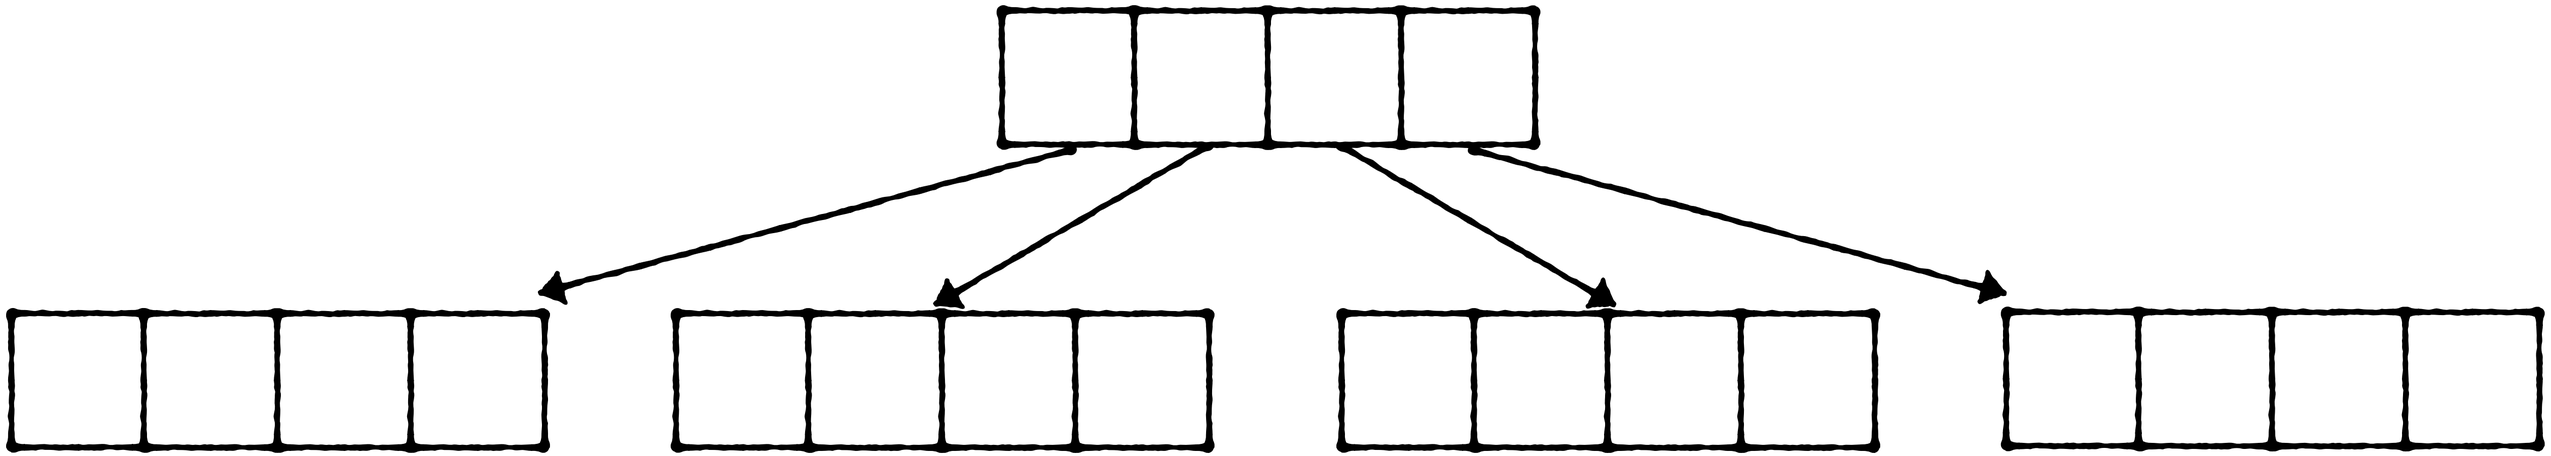

Downside: 2 memory hops per access

Mitigation: directory must fit in L1 cache

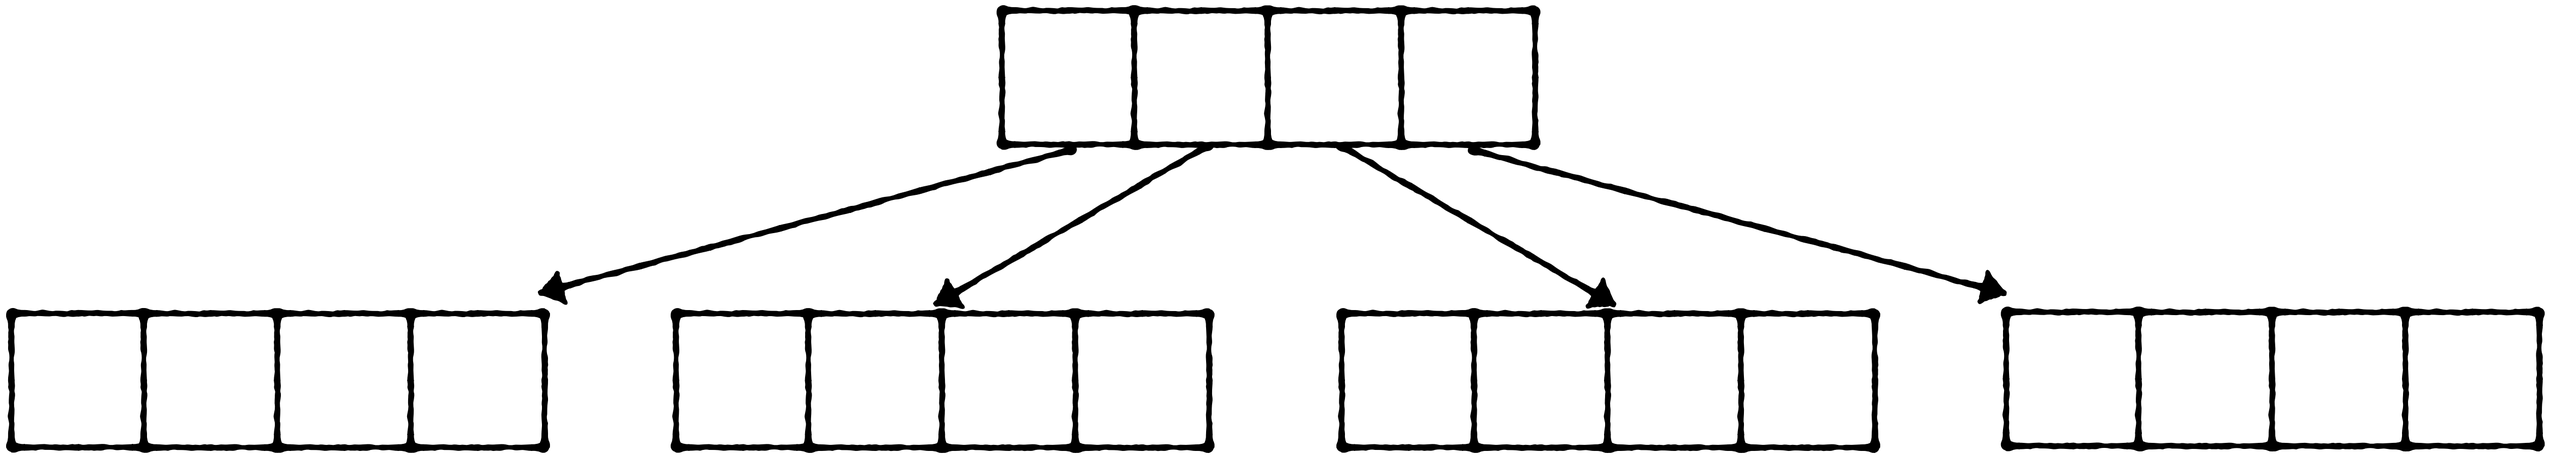

directory size?

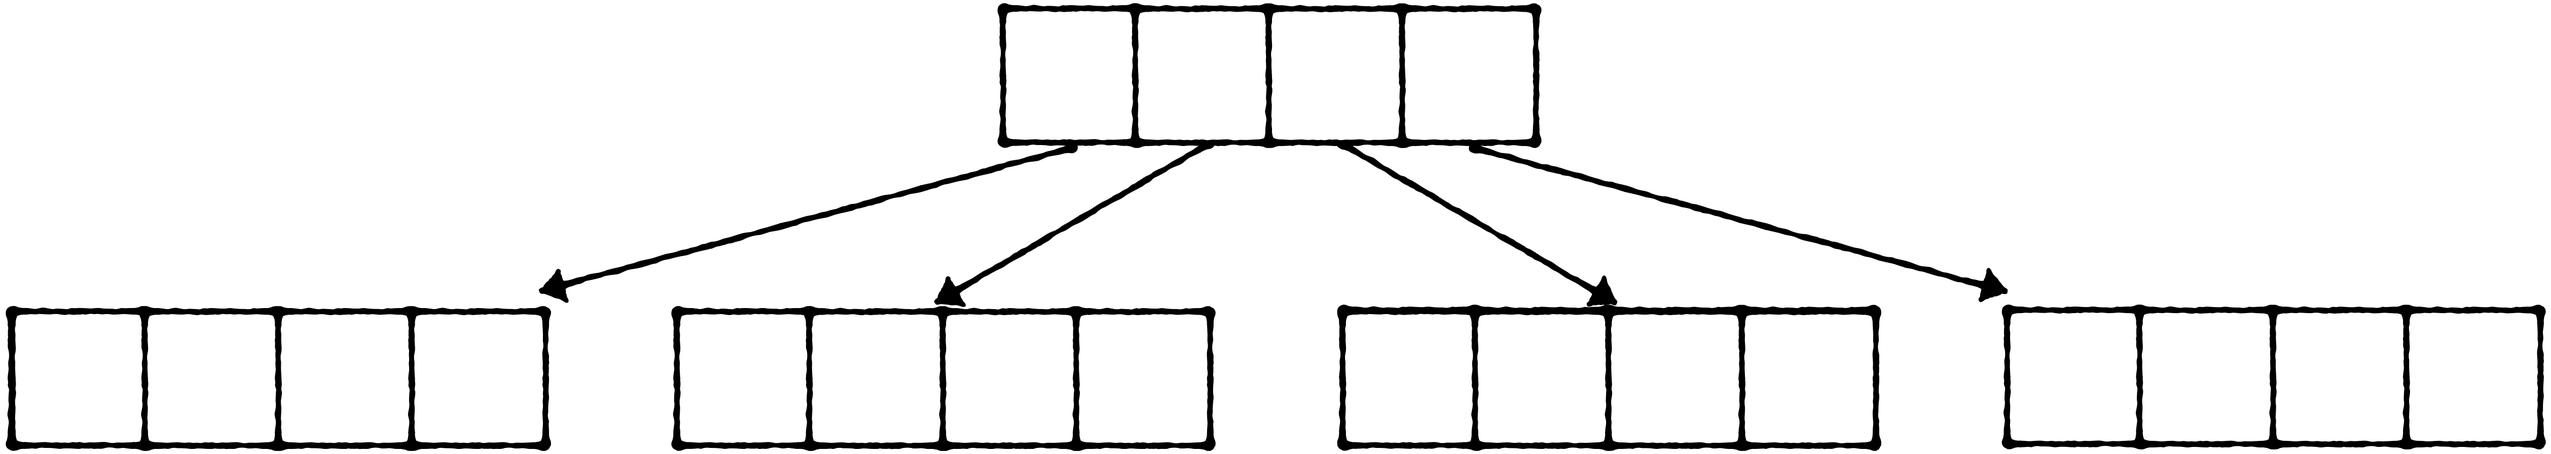

$$\text{directory size} = \frac{\text{Data size}}{\text{Data block size}}$$

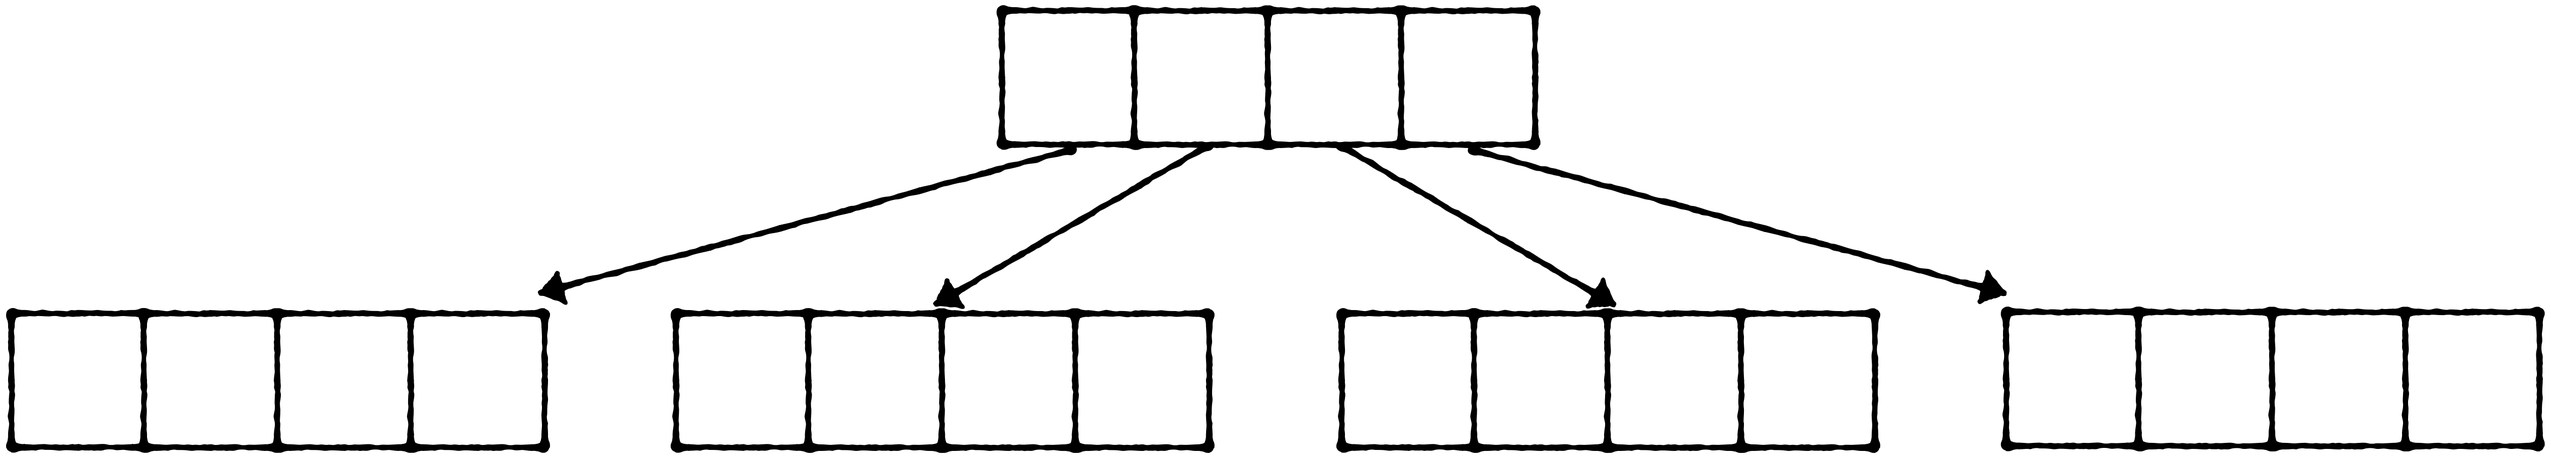

$$\text{directory size} = \frac{\text{Data size}}{\text{Data block size}} = O(N)$$

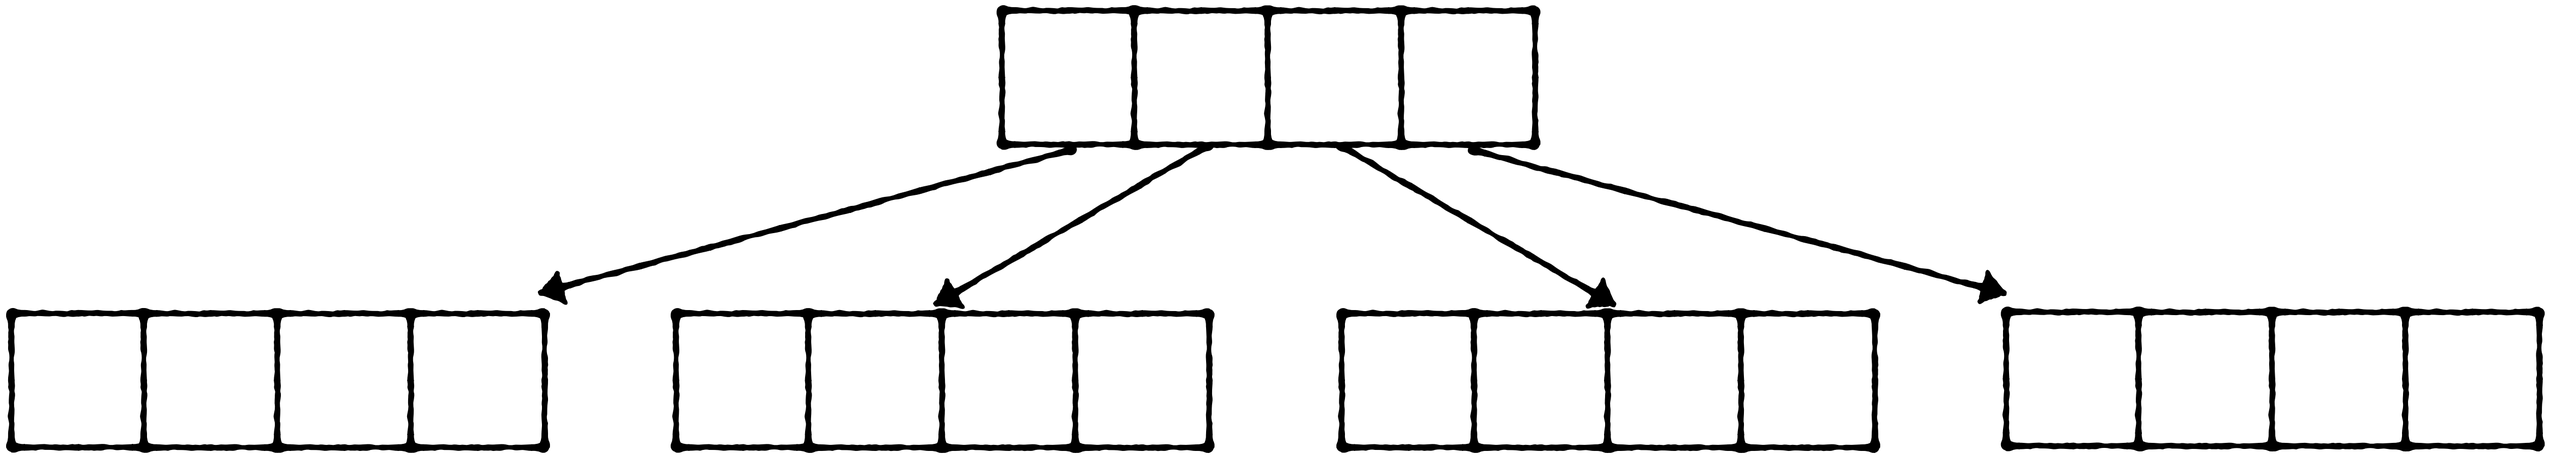

Risk: data blocks are initialized too small

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

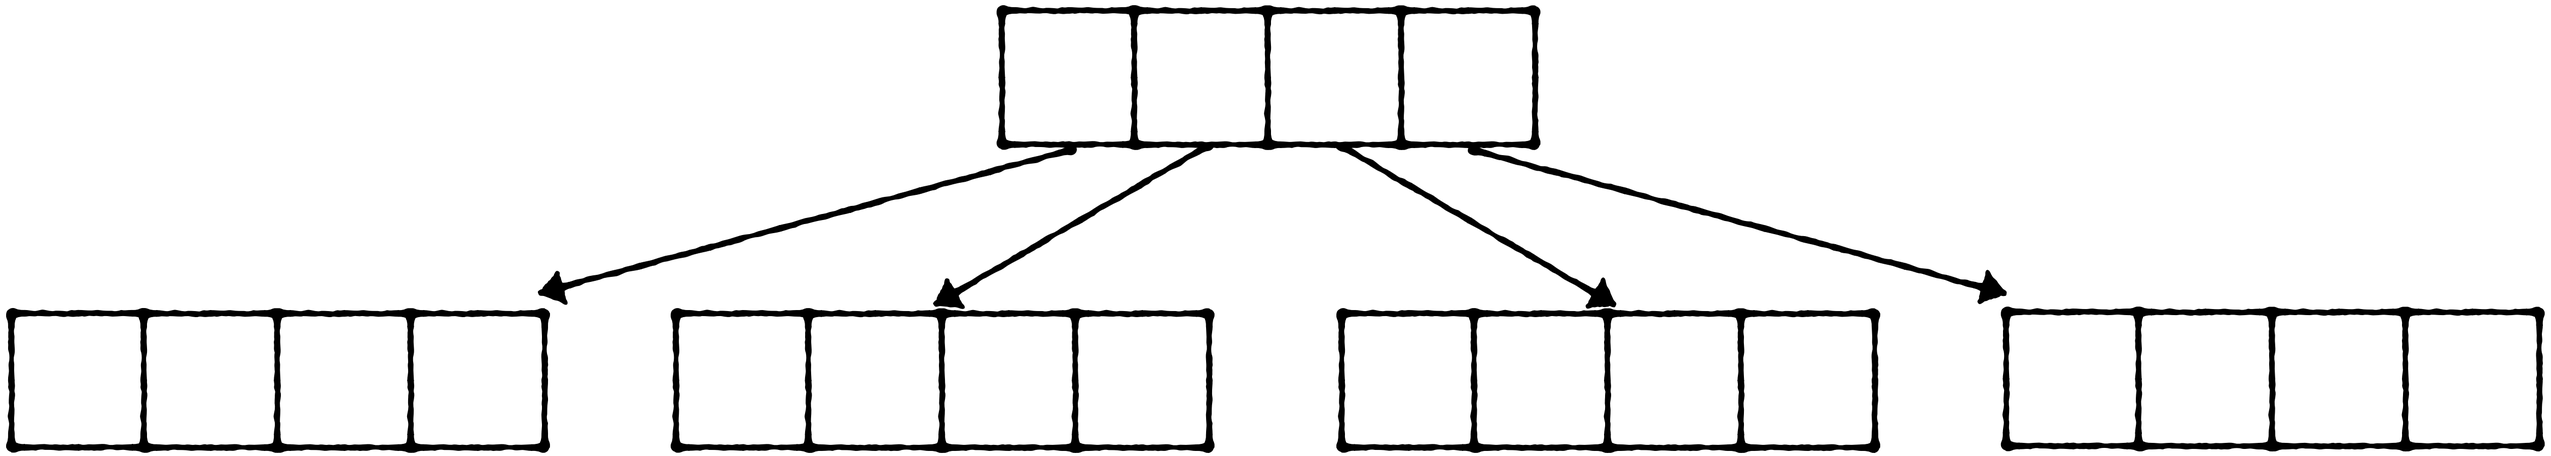

Risk: data blocks are initialized too small

Directory may outgrow the L1 cache

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

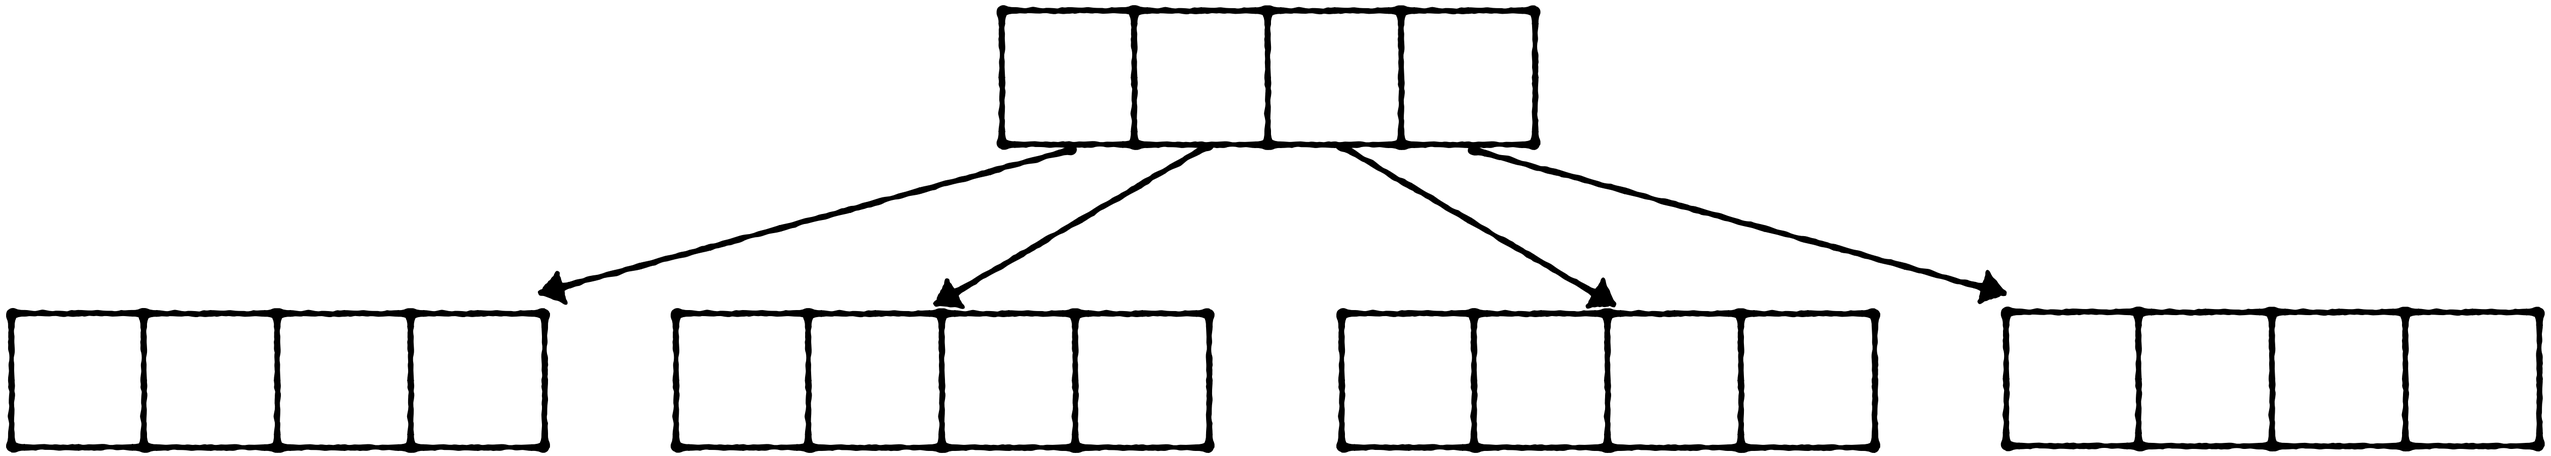

Risk: data blocks are initialized too small

Directory may outgrow the L1 cache

Solution?

Resizable Arrays in Optimal Time and Space

Algorithms and Data Structures Symposium, 1999

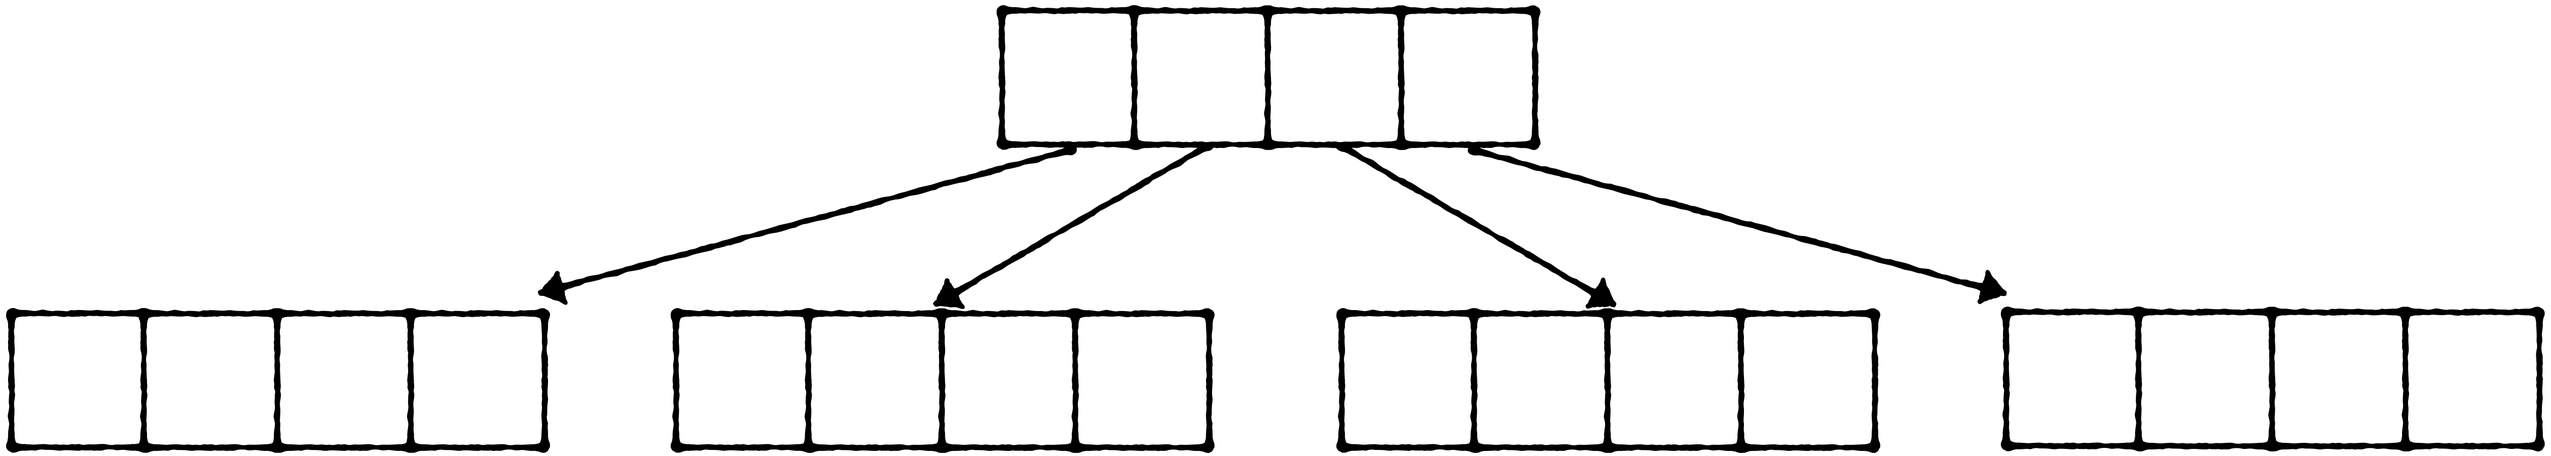

Resizable Arrays in Optimal Time and Space

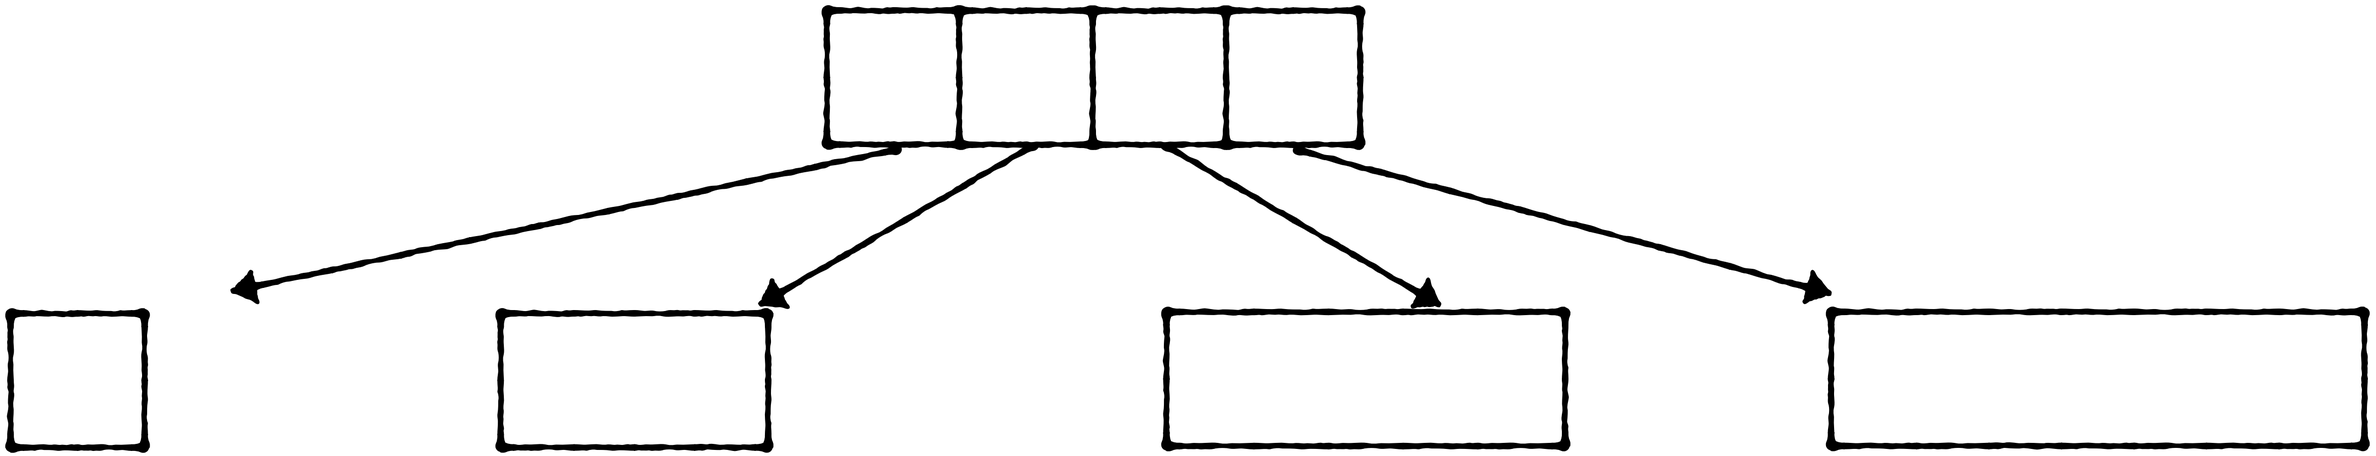

Data blocks should grow in size

Resizable Arrays in Optimal Time and Space

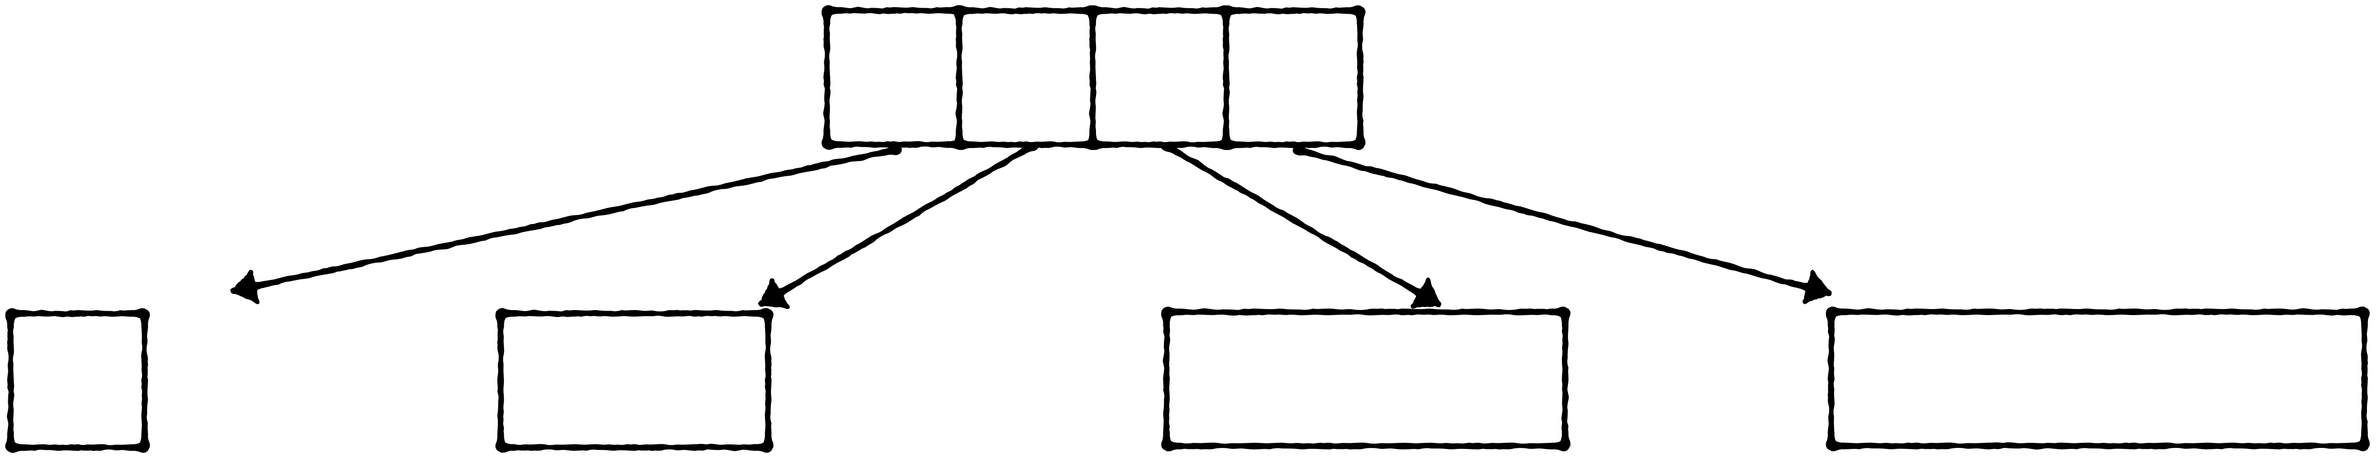

Data blocks should grow in size → **Directory grows more slowly**

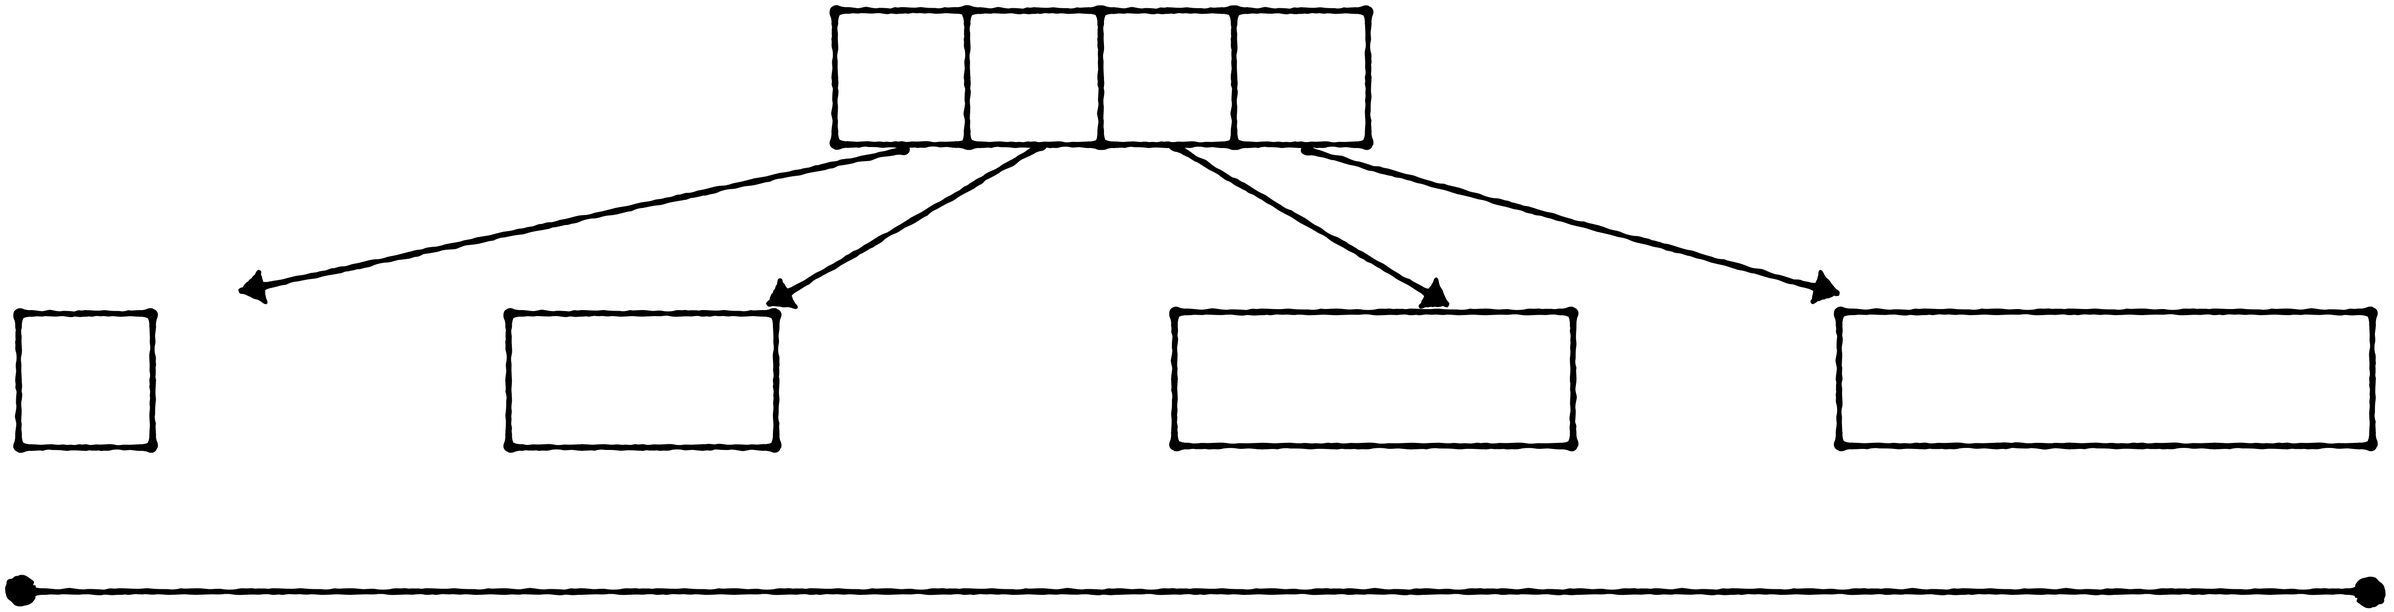

O(√N) data blocks

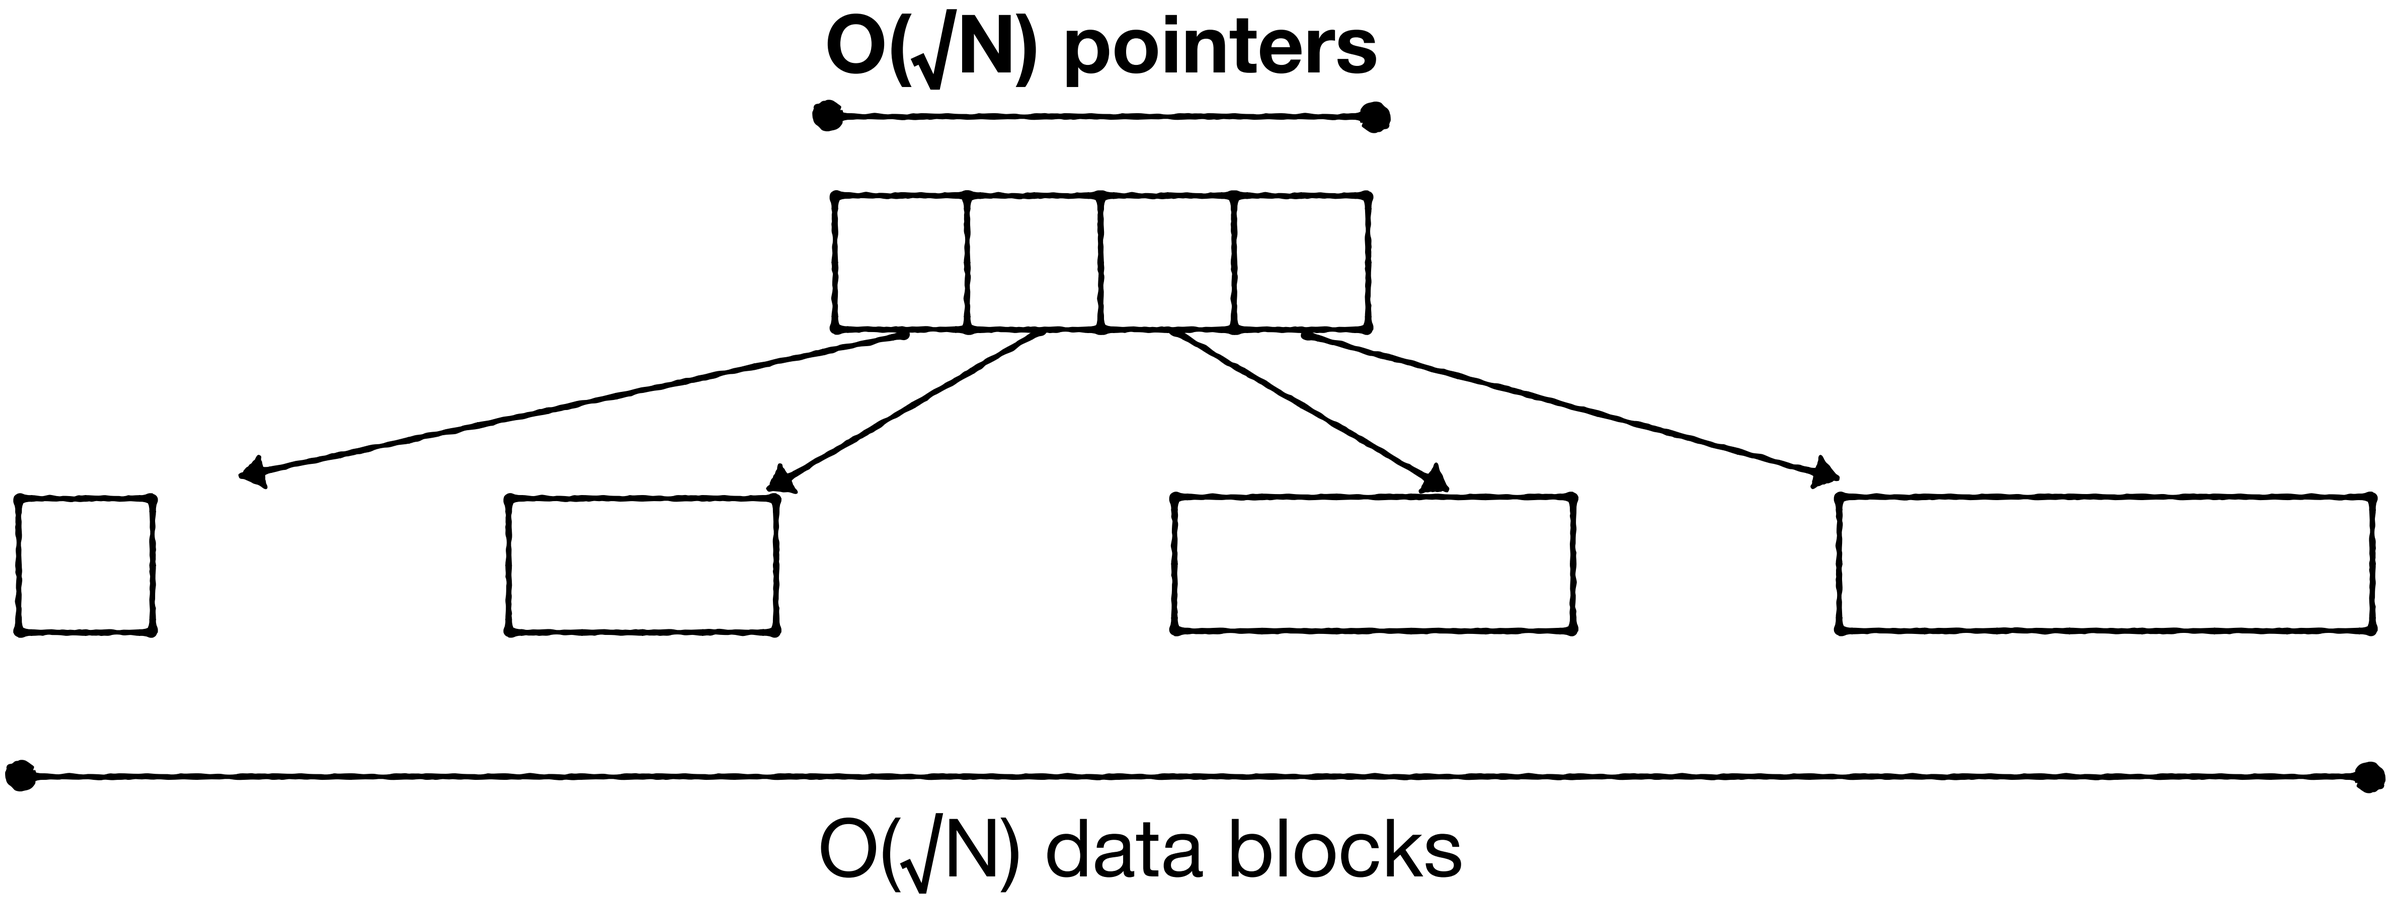
O(√N) pointers
O(√N) data blocks

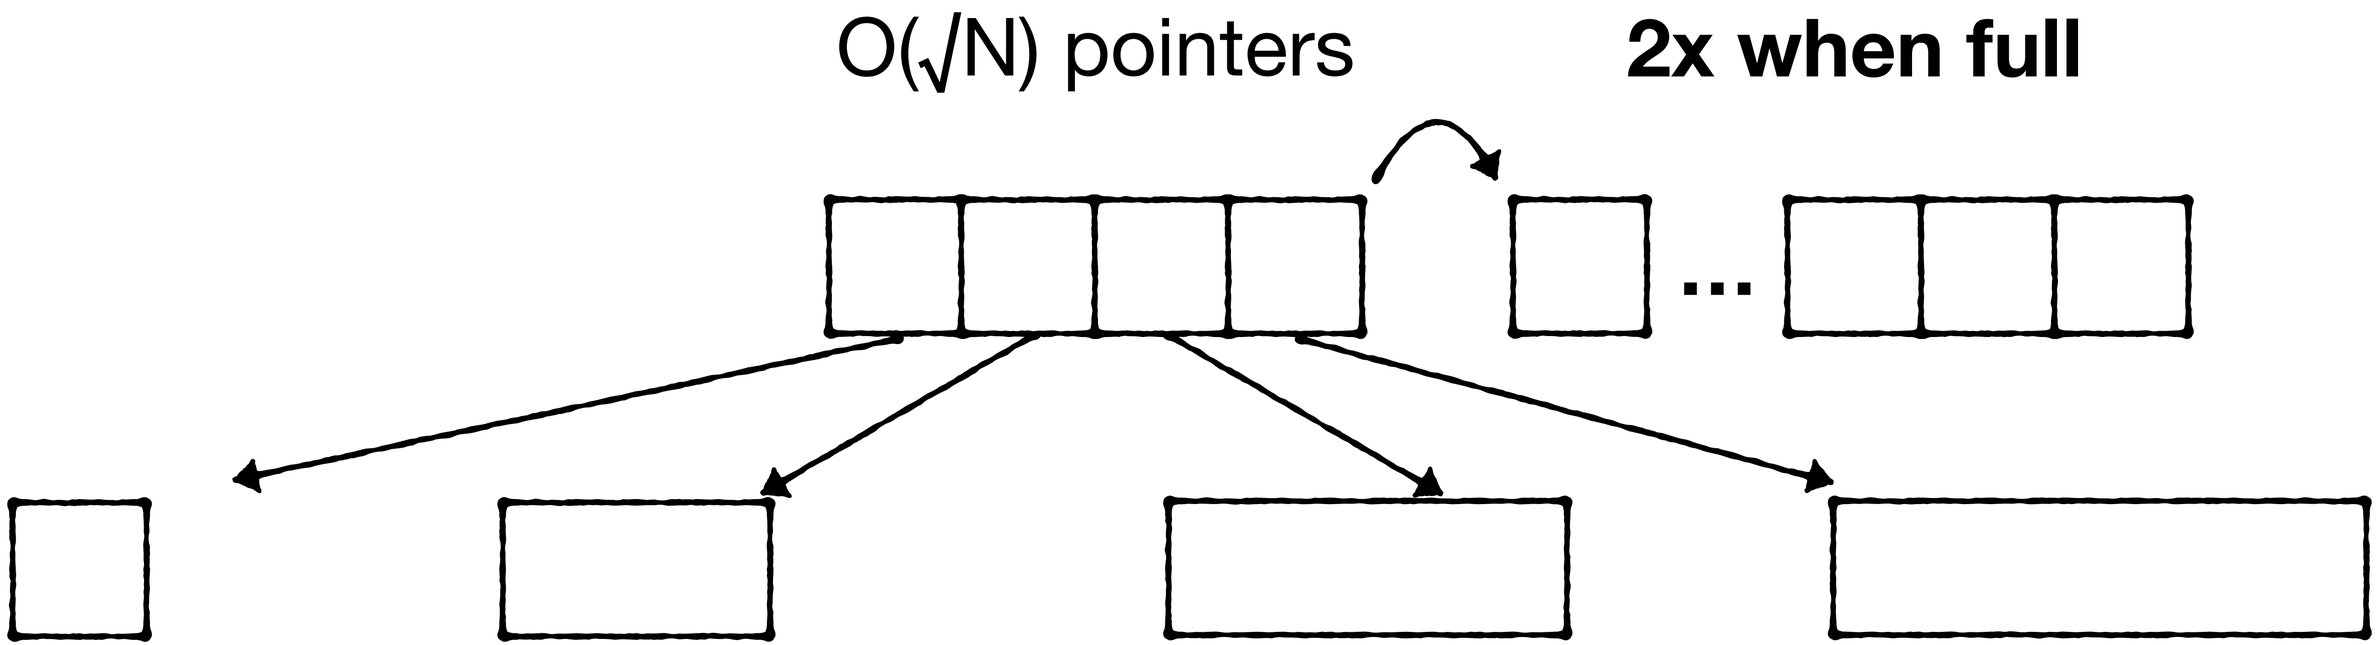
O(√N) pointers
2x when full

O(2√N) pointers

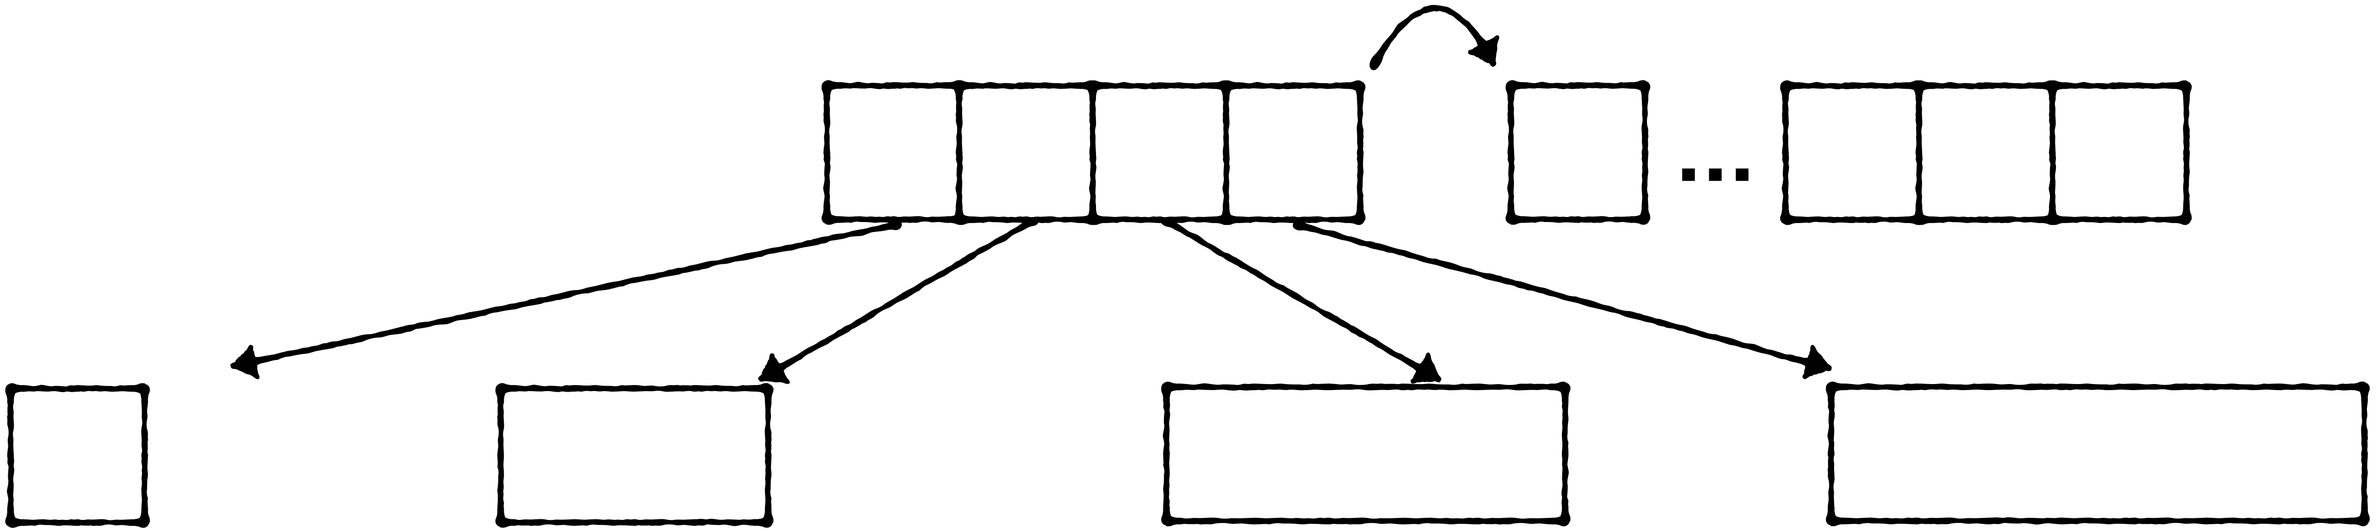

O(√N) pointers

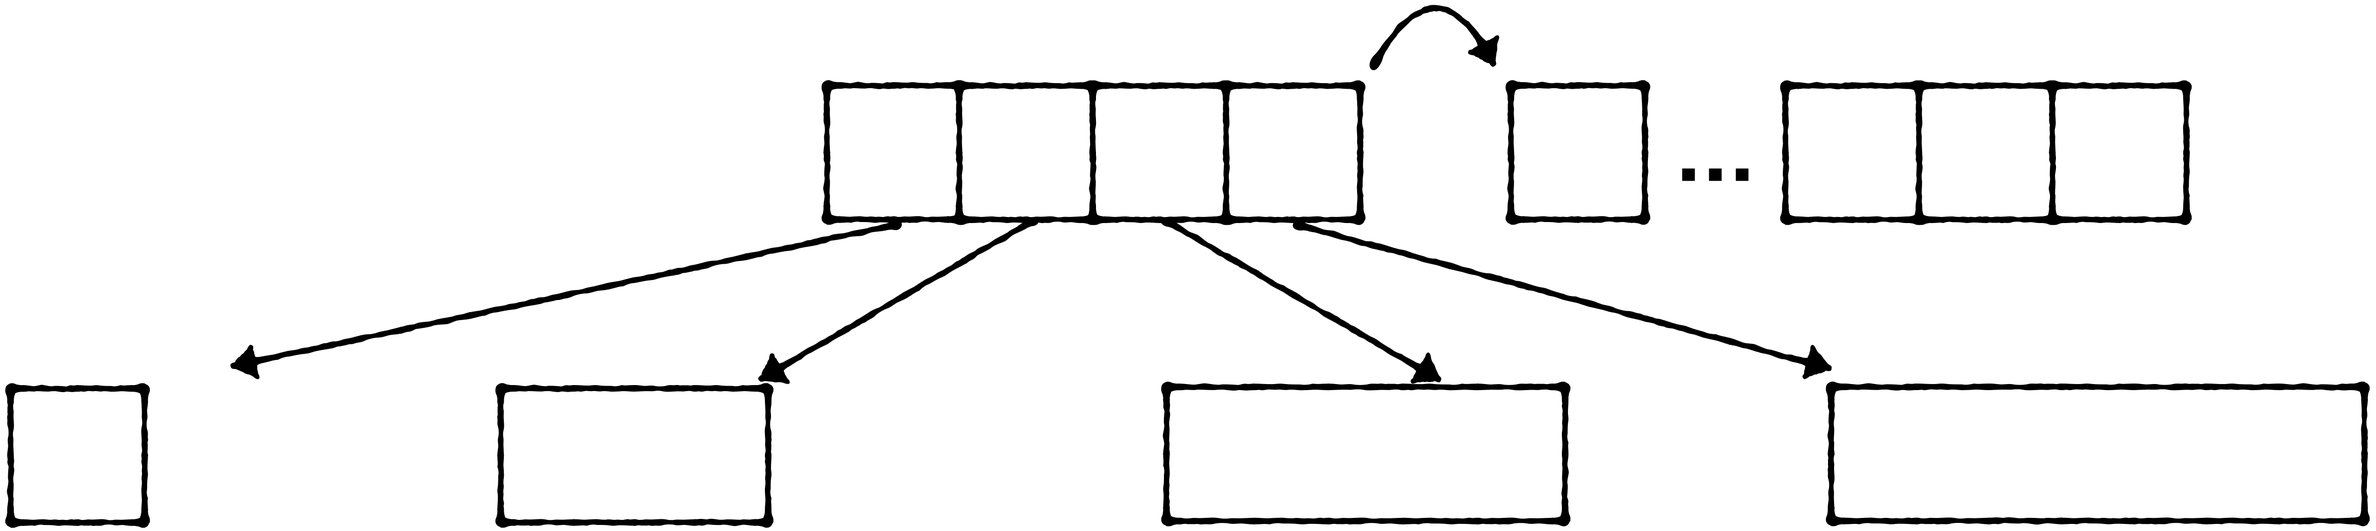

O(√N) pointers

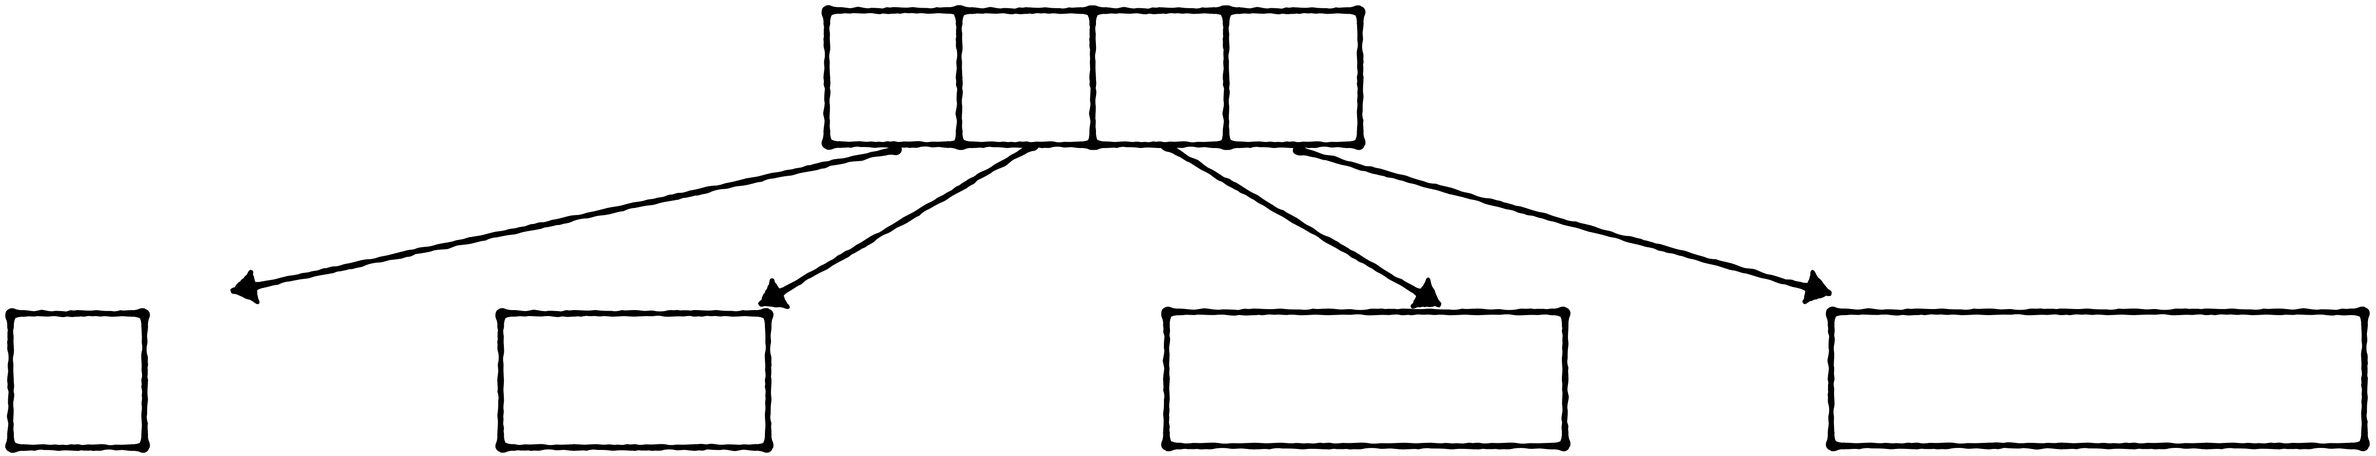

O(√N) slots

O(√N) pointers

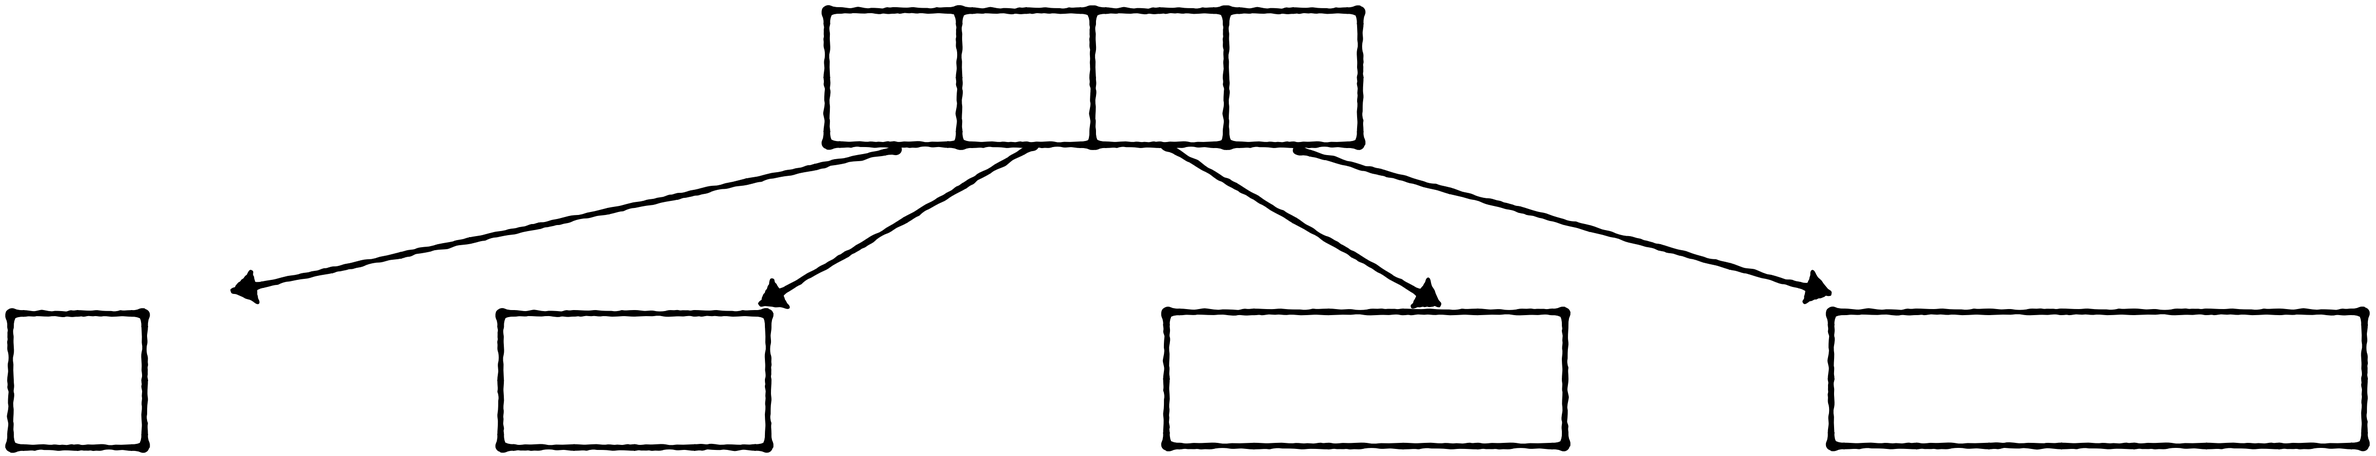

Waste at most O(√N) slots

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

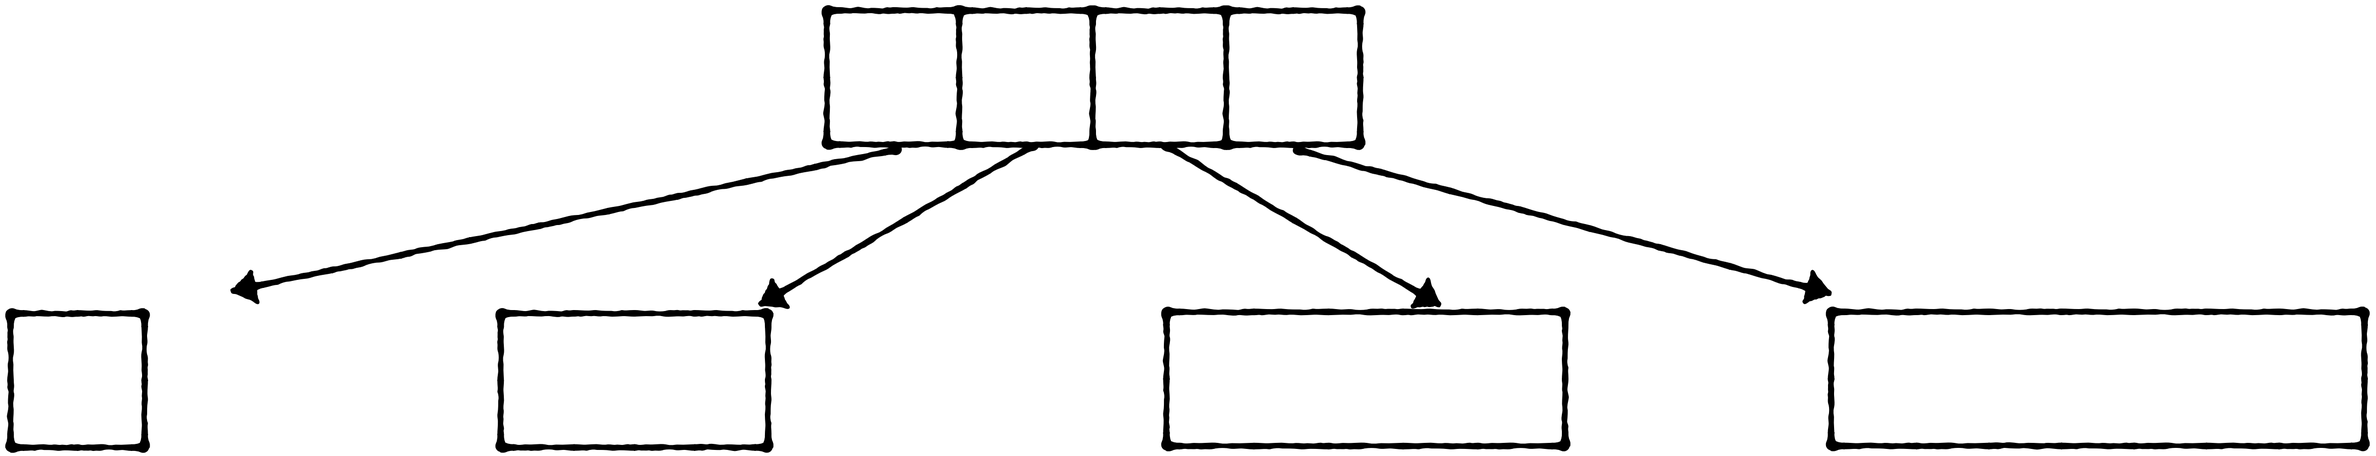

Max space amp = O(√N)

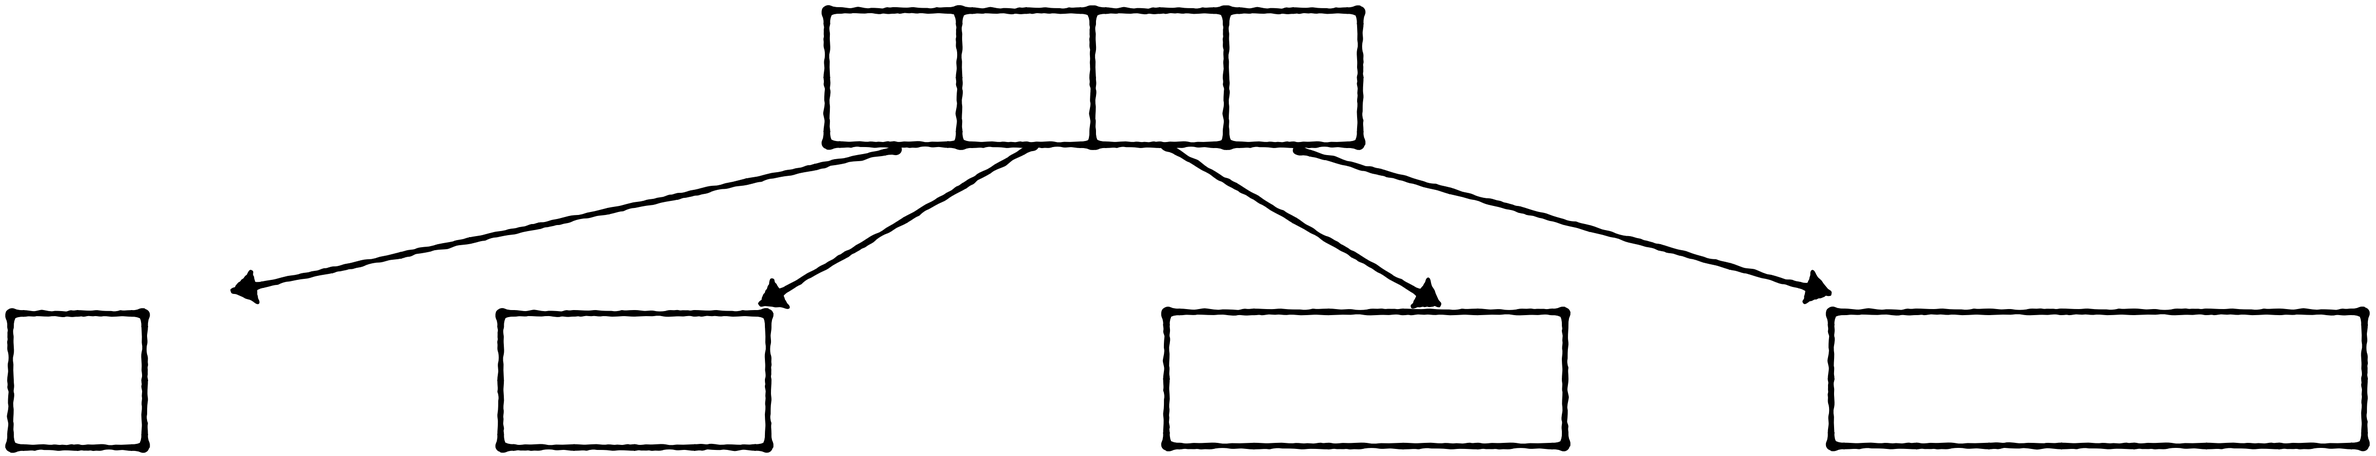

Challenges: How to grow the block to meet these properties?

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

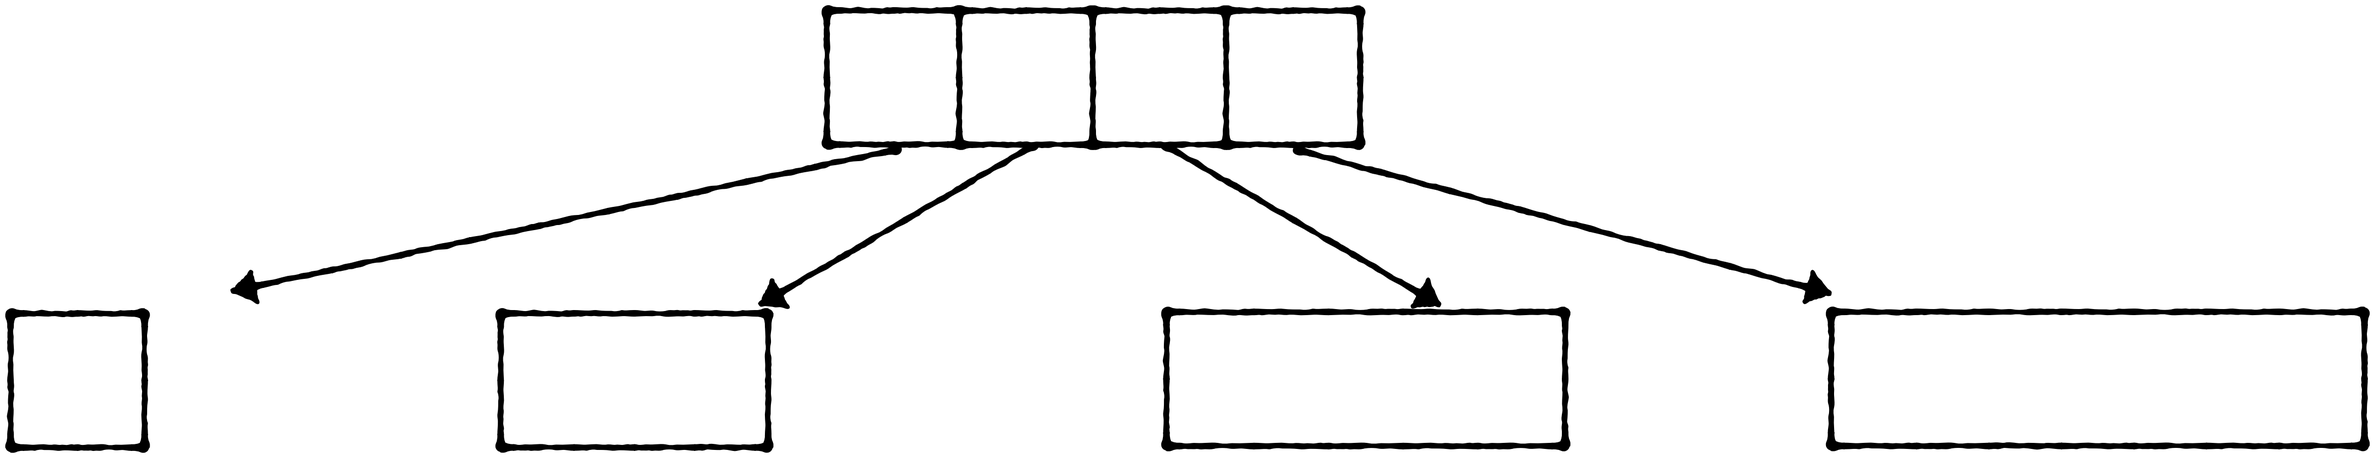

Challenges: How to grow the block to meet these properties?

Inferring which block contains which array offset?

Multiple levels

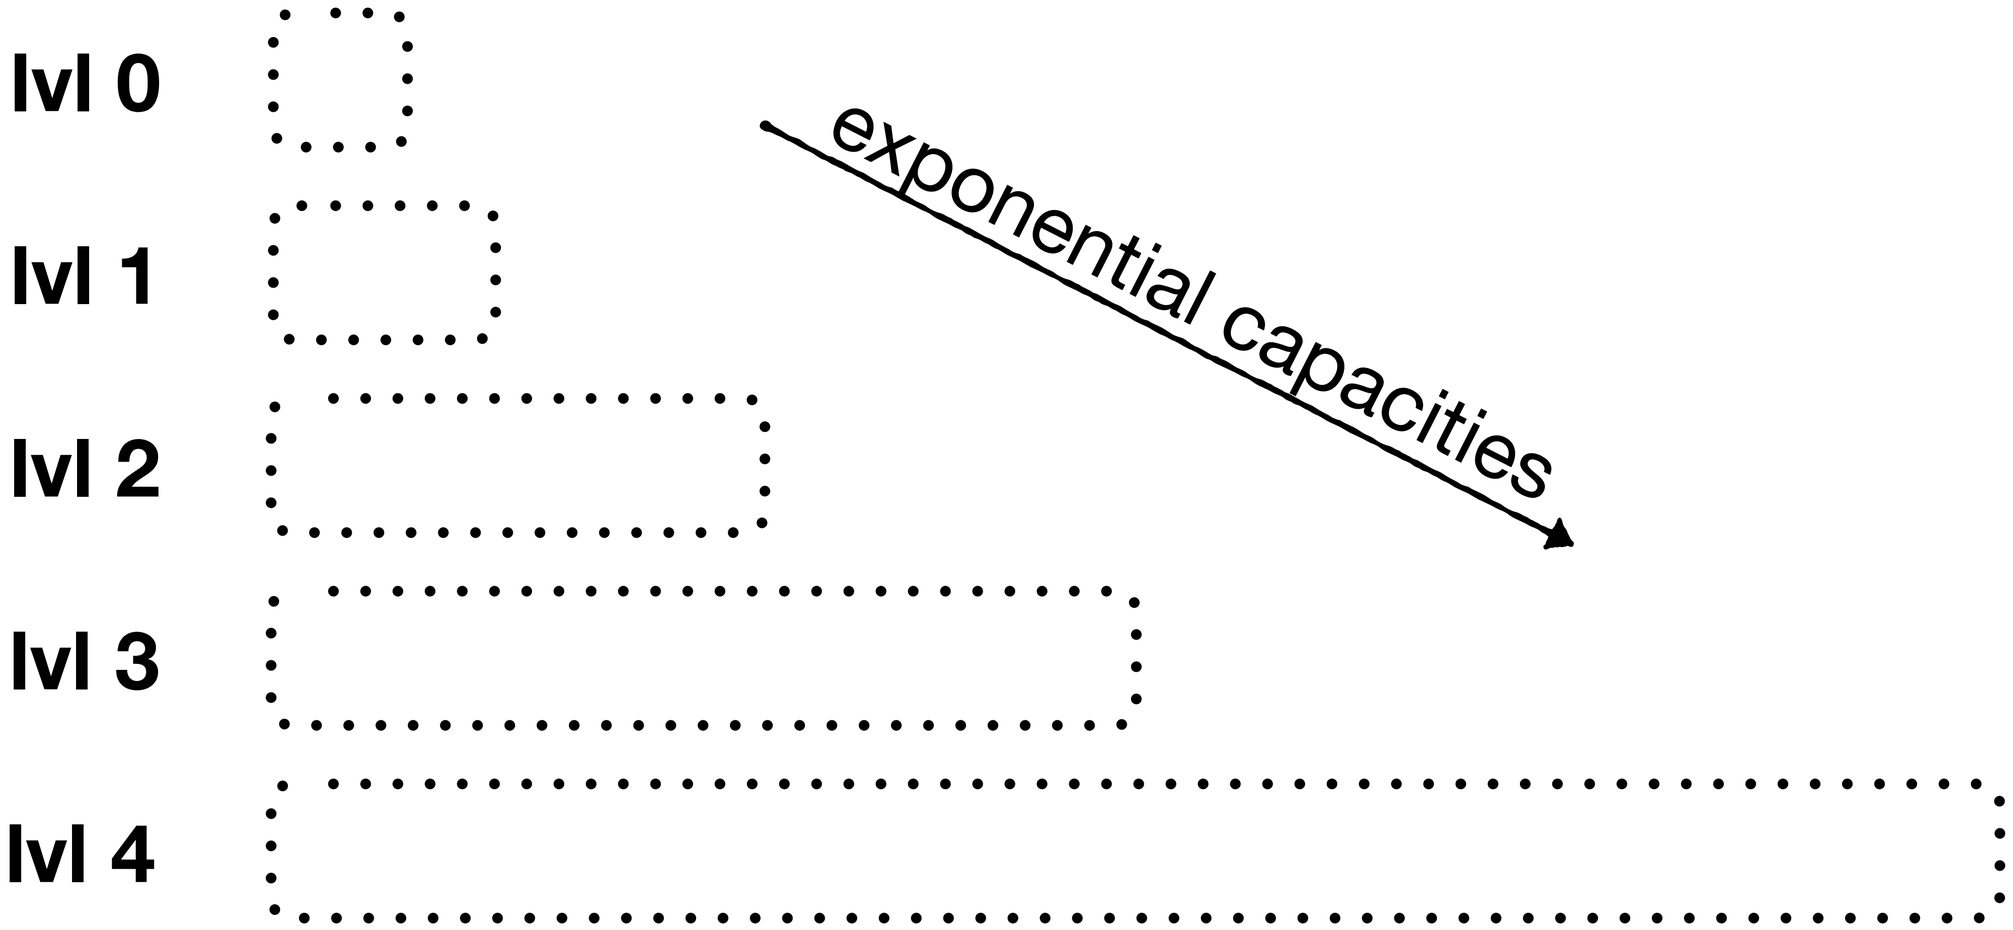

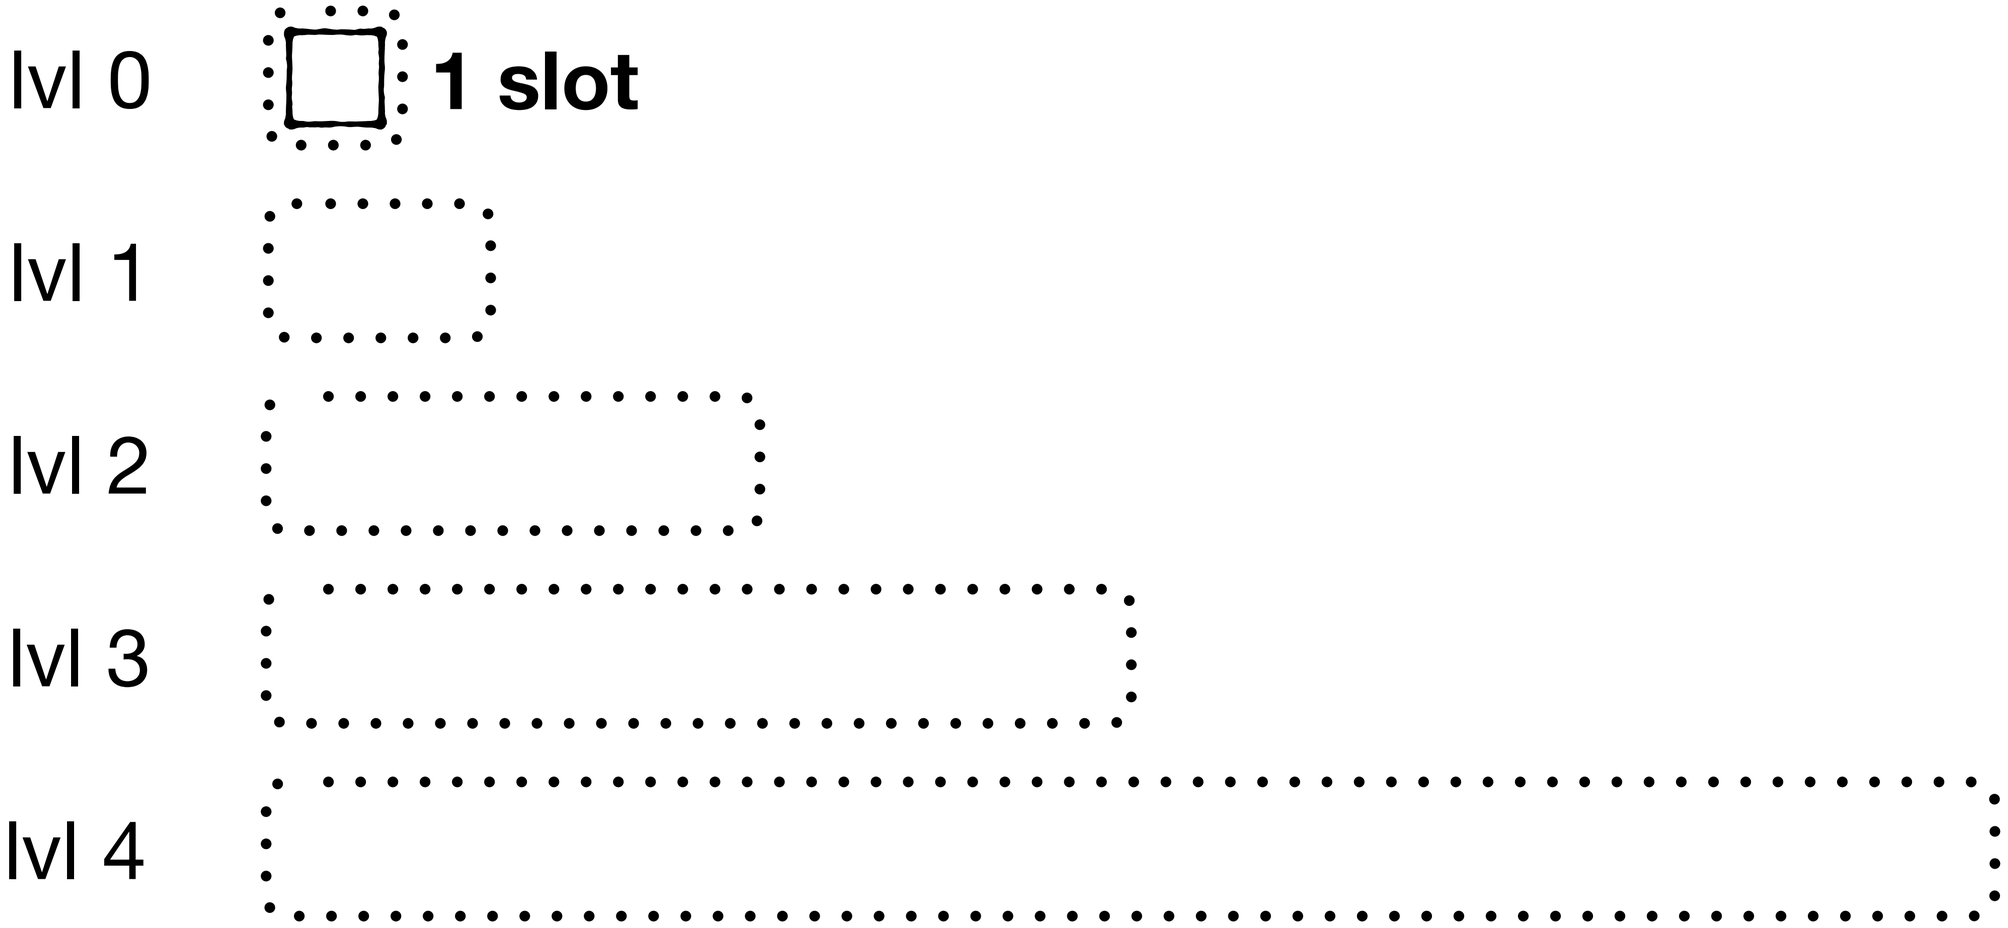
lvl 0
1 slot
lvl 1
lvl 2
lvl 3
lvl 4

In every pair of subsequent levels k and k+1

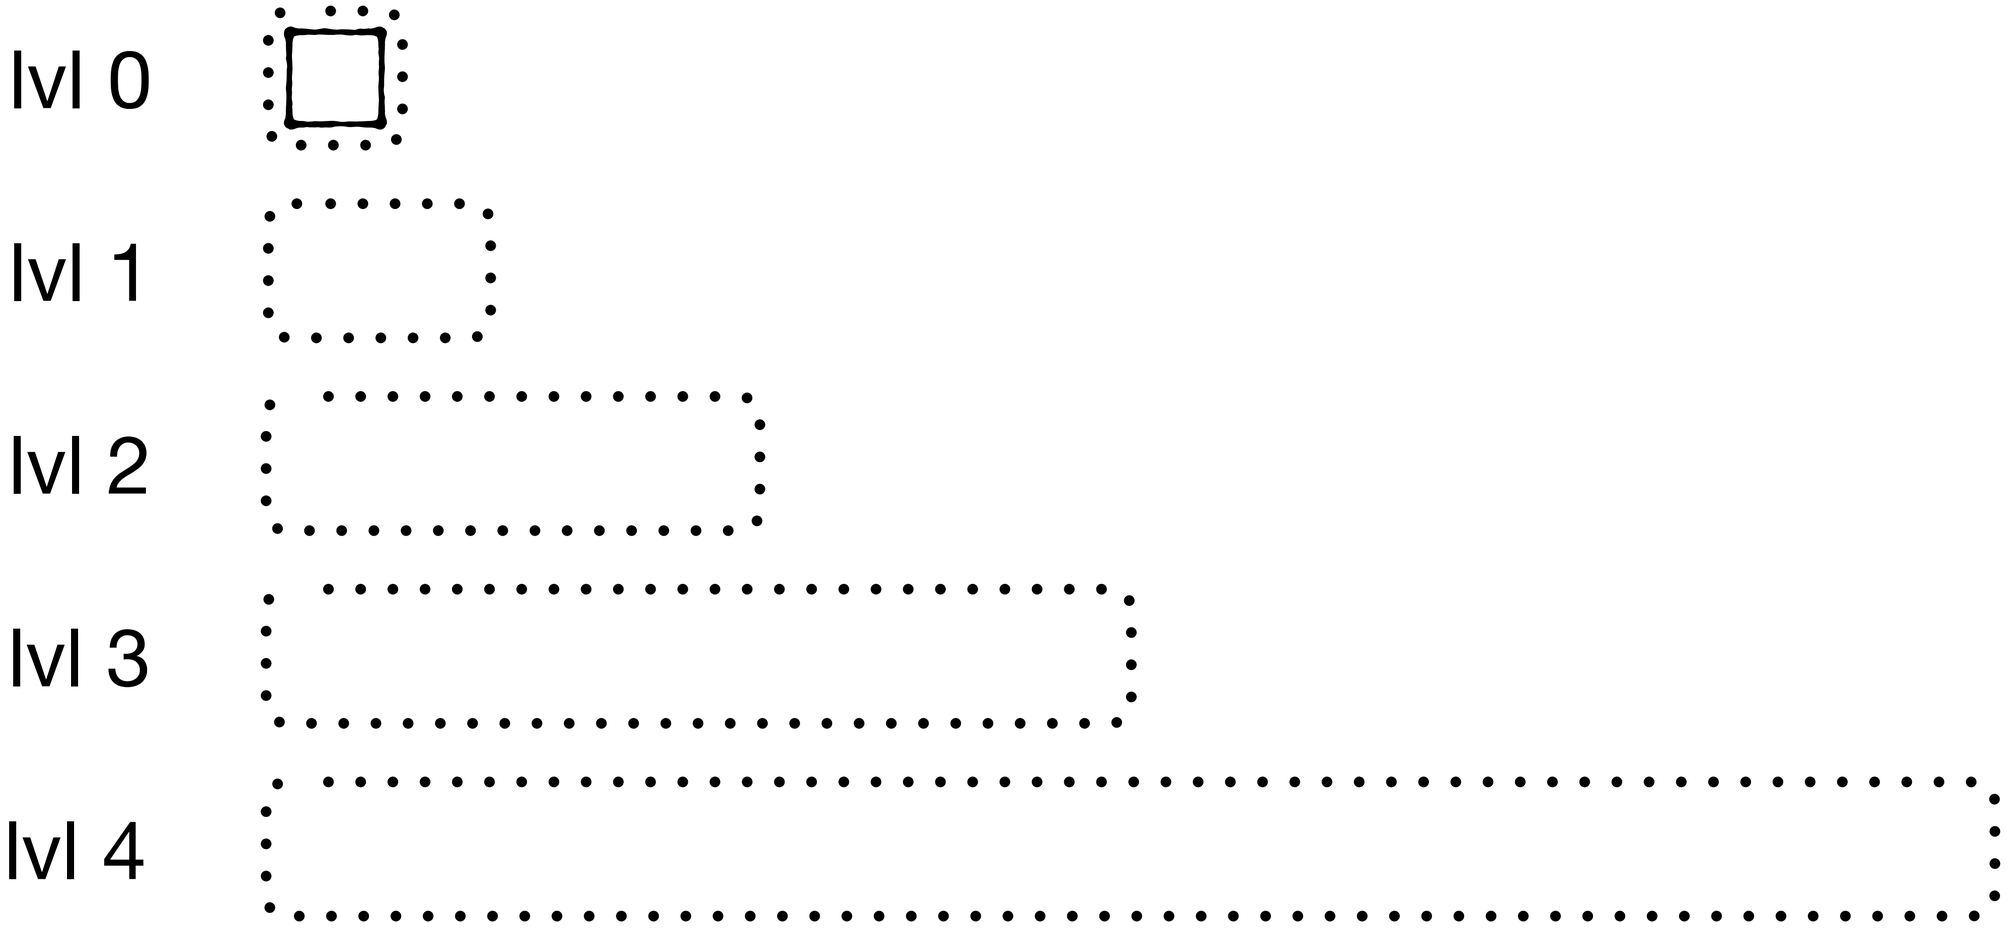

In every pair of subsequent levels k and k+1

Size of arrays doubles at level k

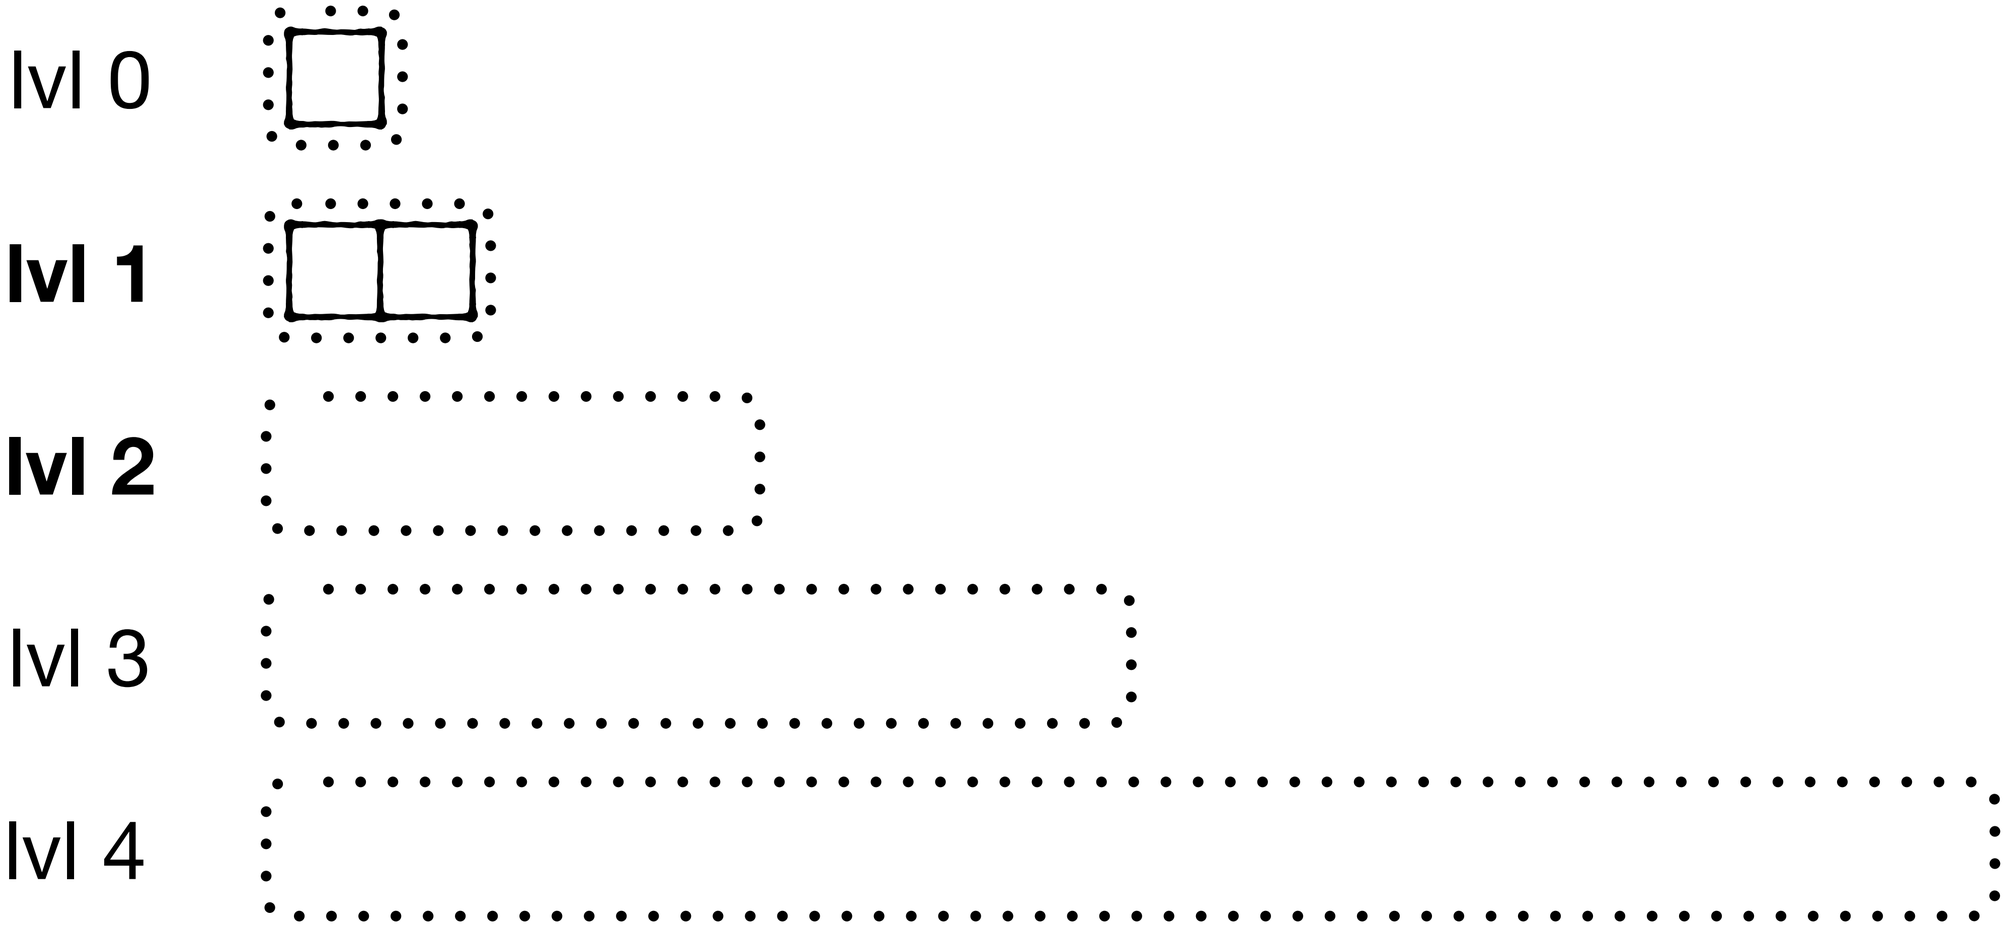

In every pair of subsequent levels k and k+1

Size of arrays doubles at level k

 # arrays doubles at level k+1

lvl 0

lvl 1

lvl 2

lvl 3

lvl 4

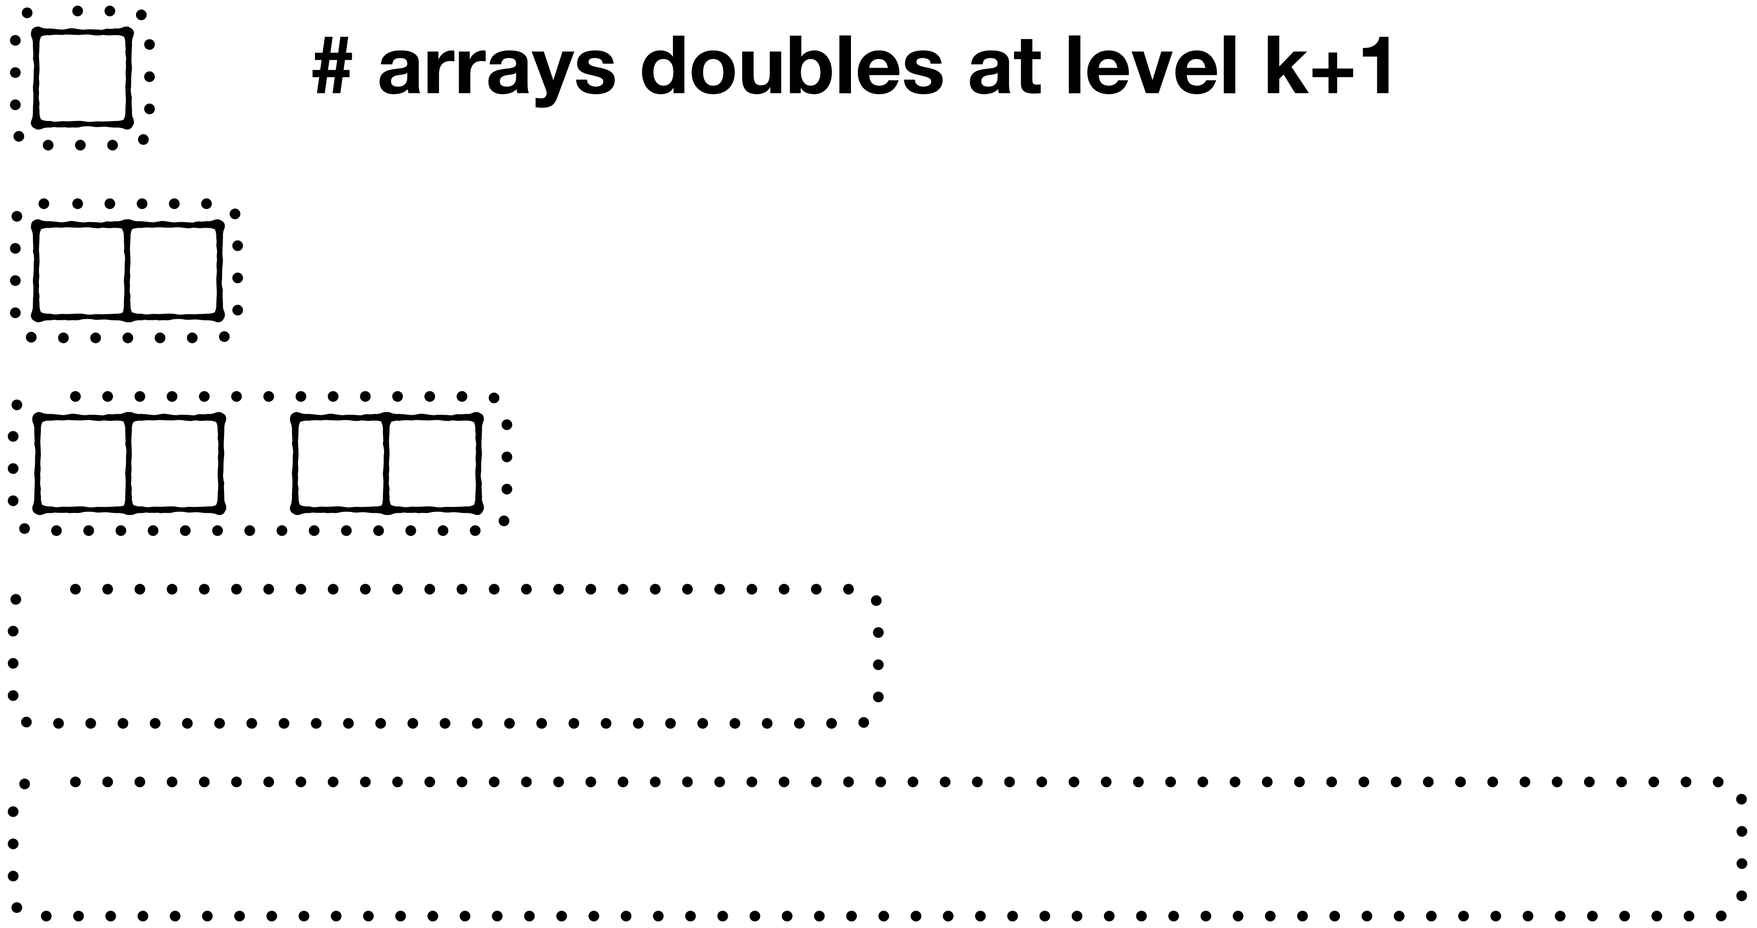

In every pair of subsequent levels k and k+1

Size of arrays doubles at level k

arrays doubles at level k+1

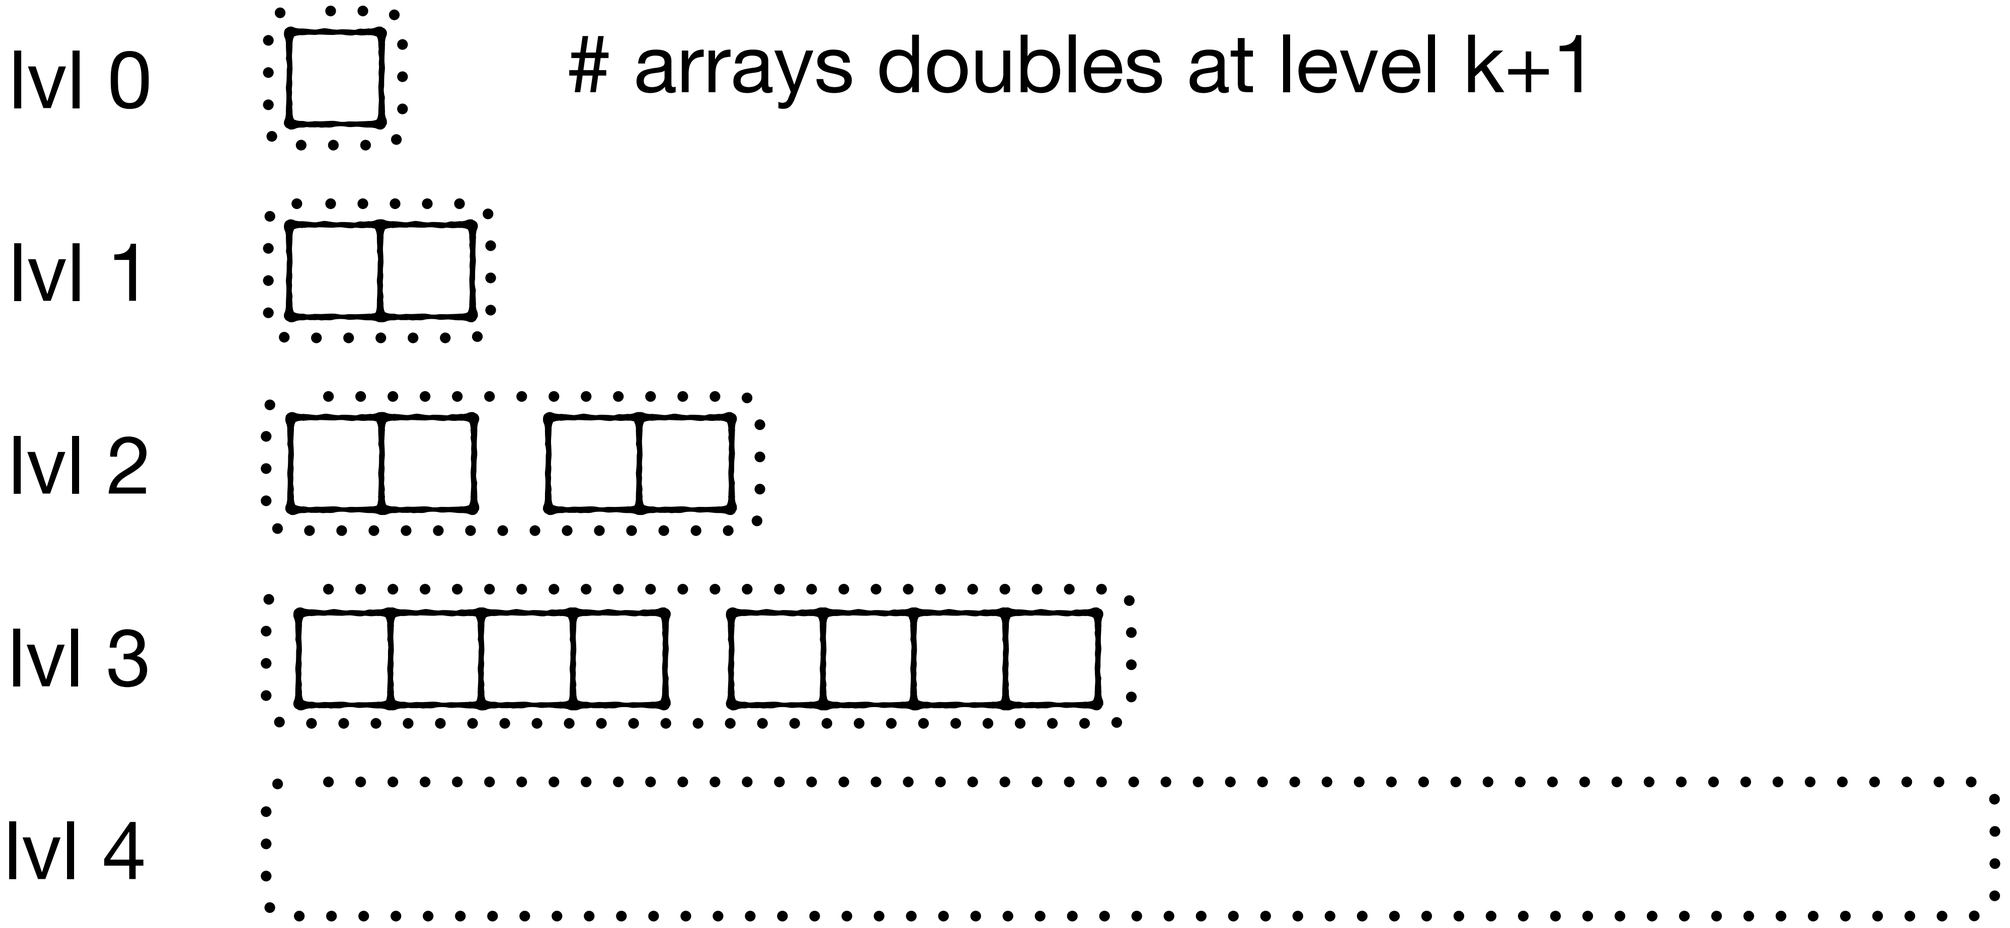

In every pair of subsequent levels k and k+1

Size of arrays doubles at level k

 # arrays doubles at level k+1

lvl 0

lvl 1

lvl 2

lvl 3

lvl 4

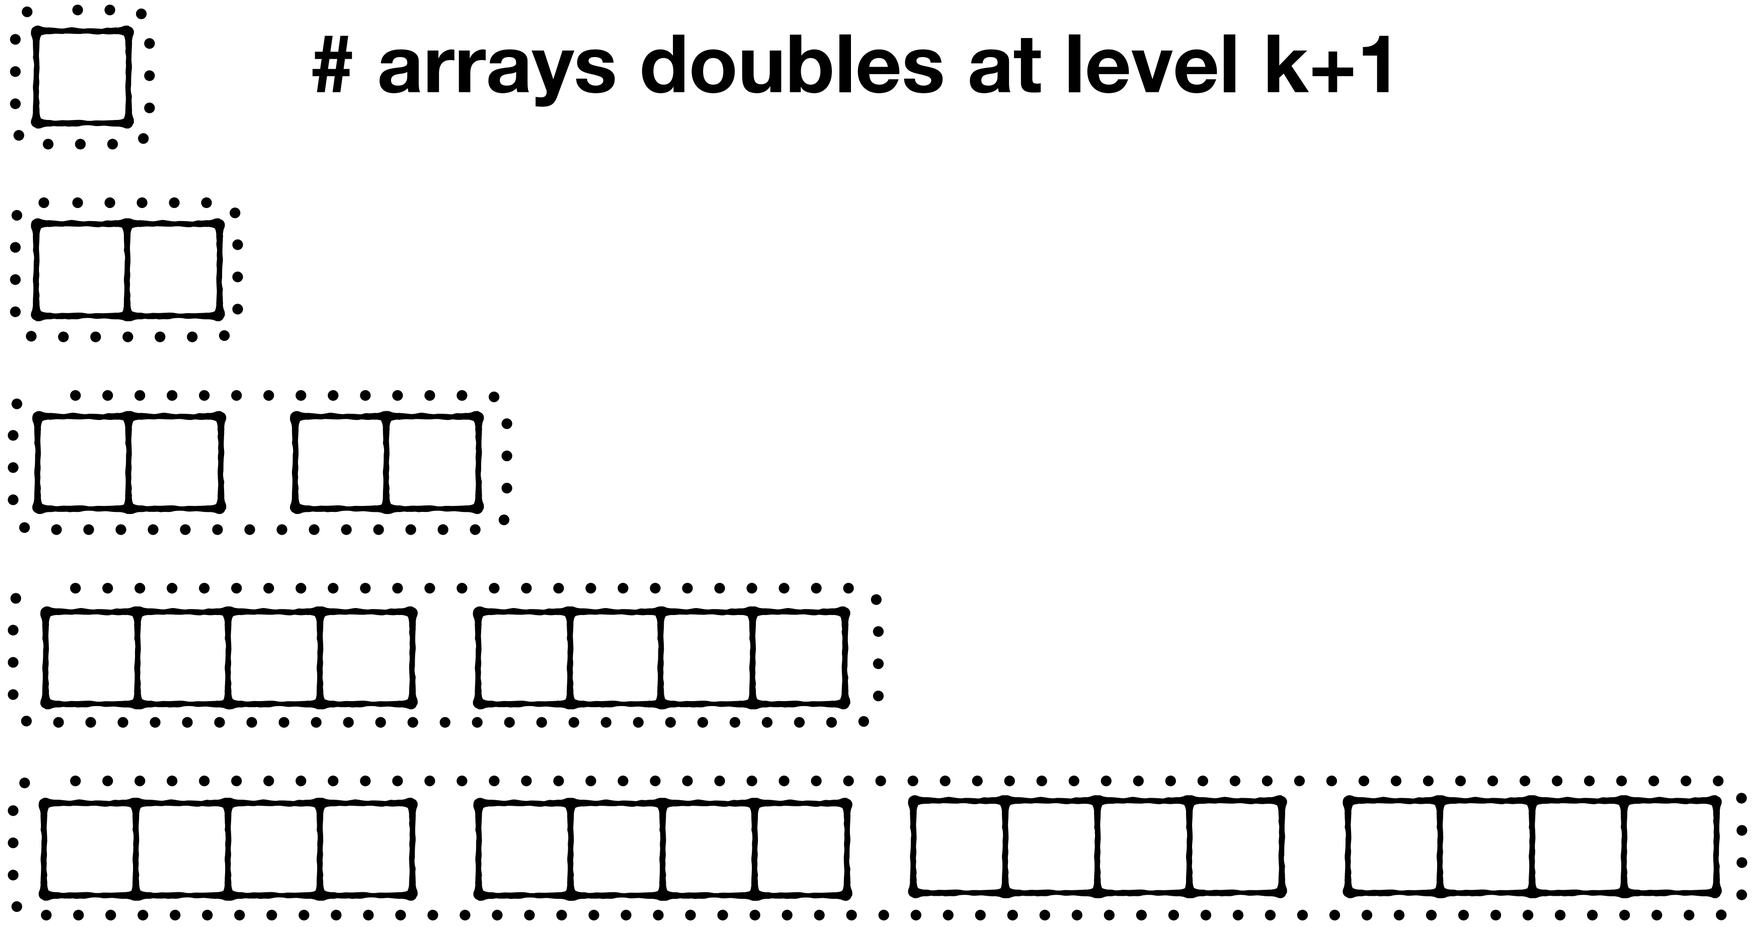

lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

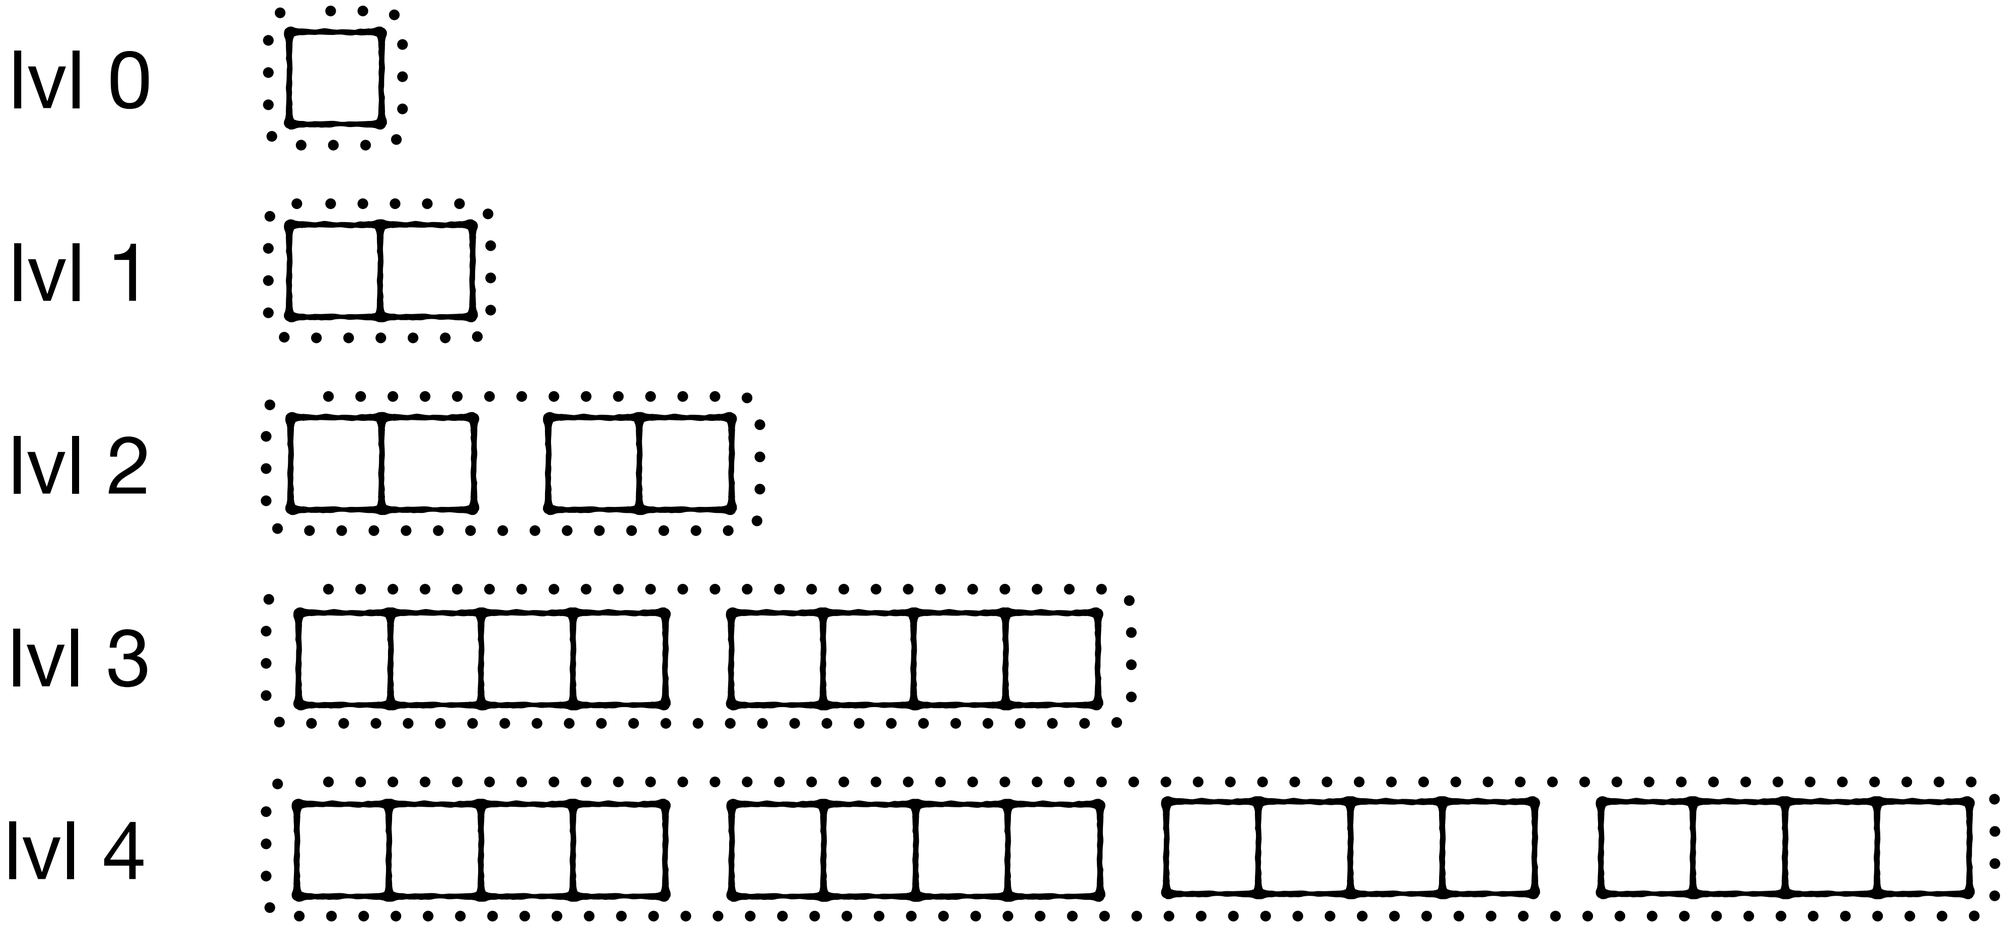

lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

lvl 0 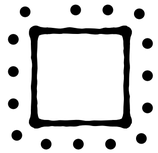 $2^{\lfloor 0/2 \rfloor} = 1$

lvl 1 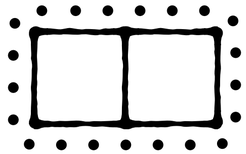 $2^{\lfloor 1/2 \rfloor} = 1$

lvl 2 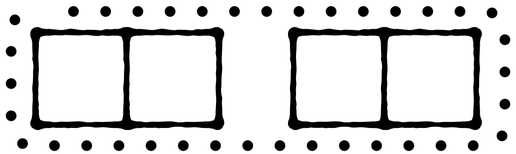 $2^{\lfloor 2/2 \rfloor} = 2$

lvl 3 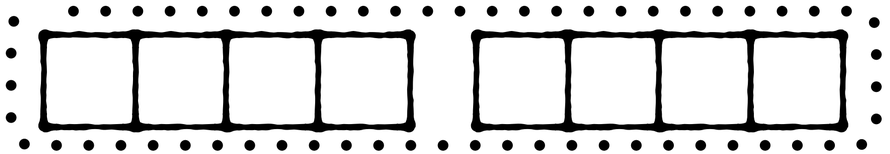 $2^{\lfloor 3/2 \rfloor} = 2$

$2^{\lfloor 4/2 \rfloor} = 4$

lvl 4 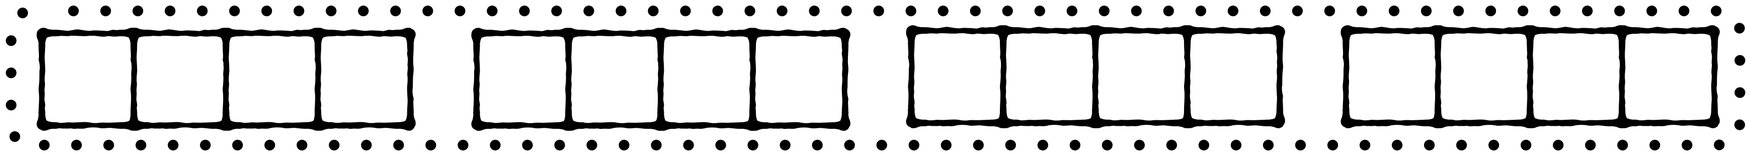

lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, **each with $2^{\lceil K/2 \rceil}$ slots**

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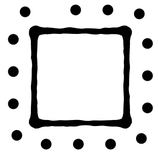 $2^{\lceil 0/2 \rceil} = 1$

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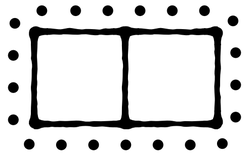 $2^{\lceil 1/2 \rceil} = 2$

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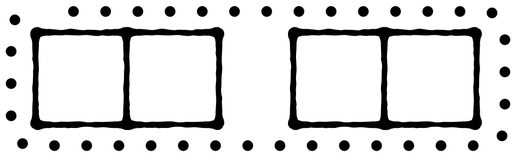 $2^{\lceil 2/2 \rceil} = 2$

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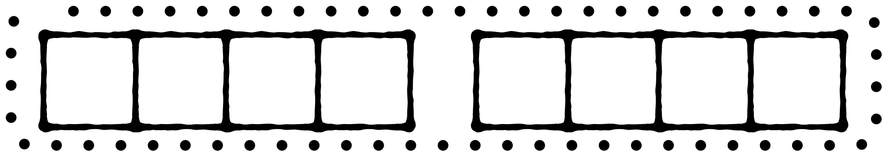 $2^{\lceil 3/2 \rceil} = 4$

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

lvl 4 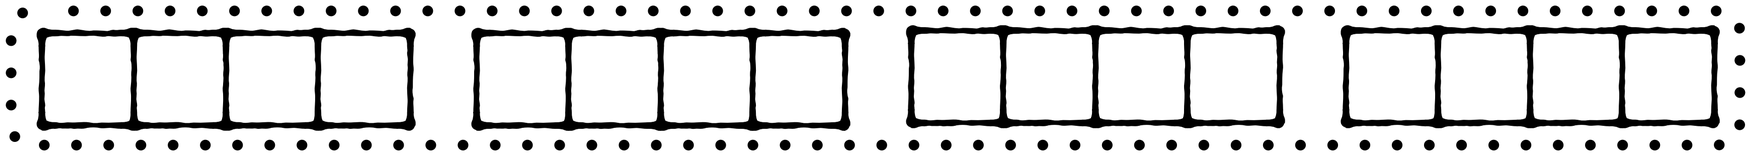

lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

 # levels?

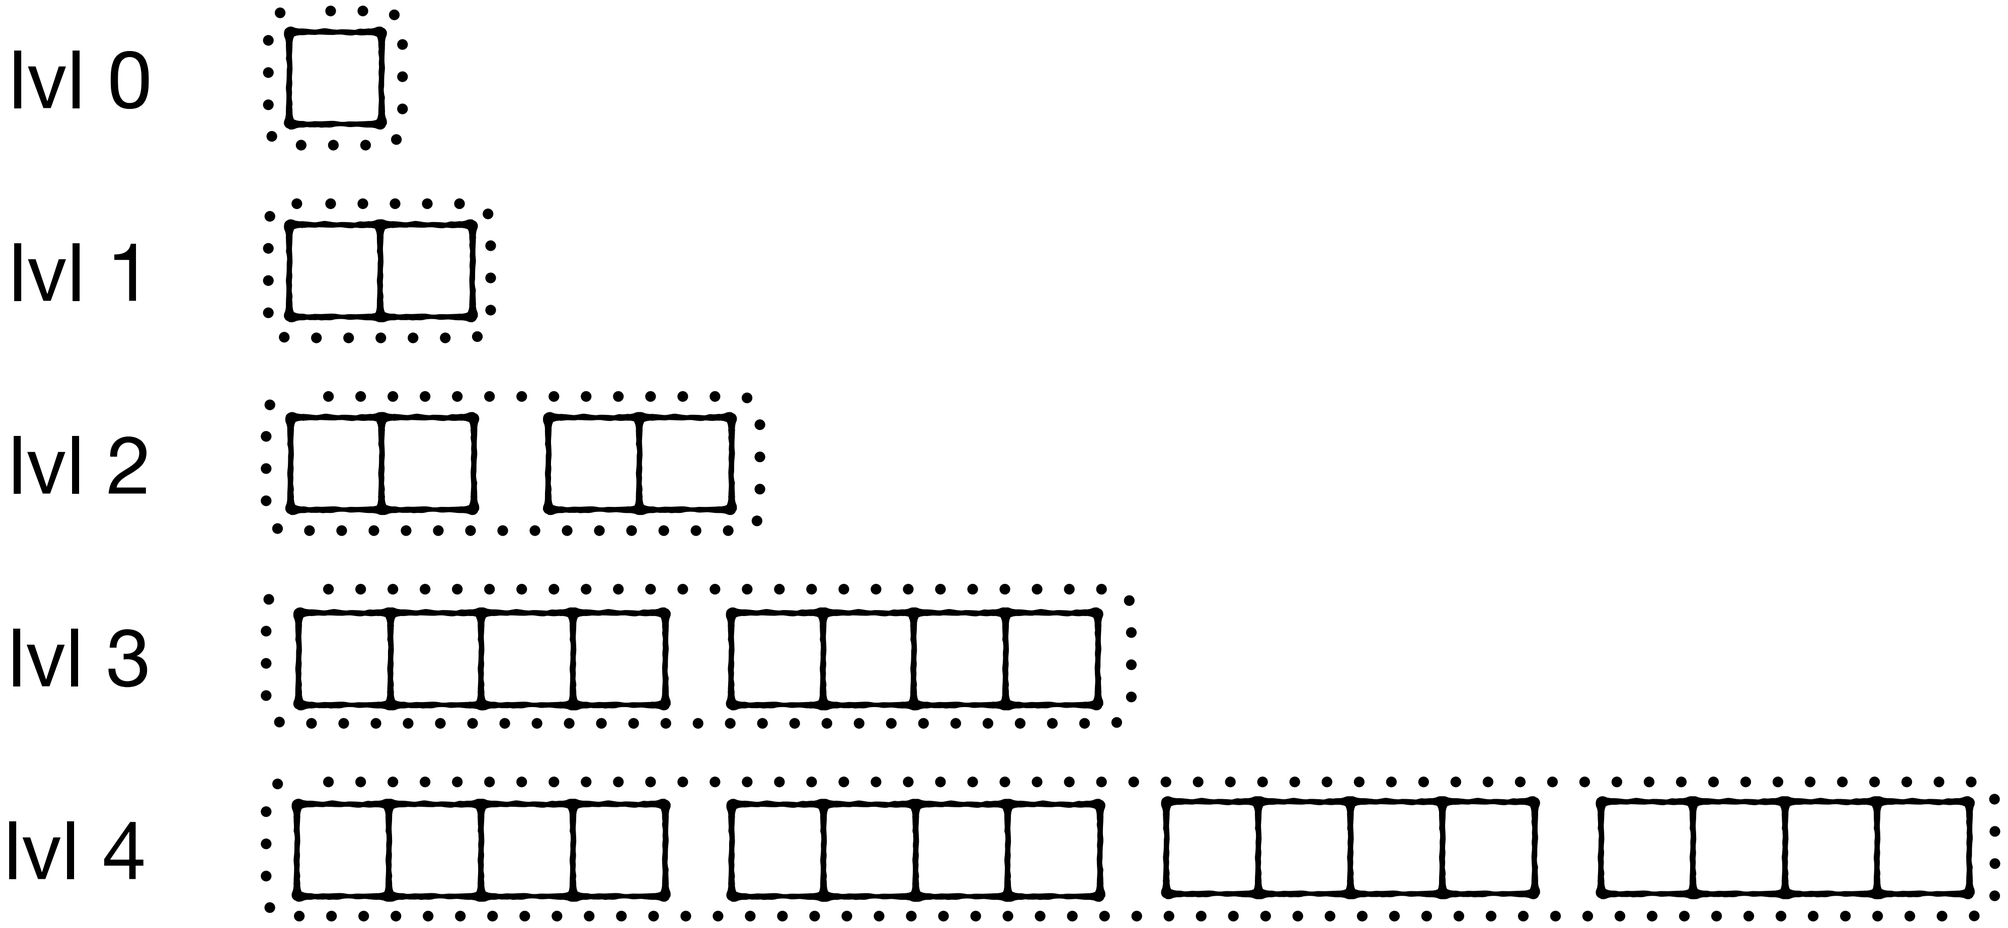

lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

 # levels: $\log_2 N$

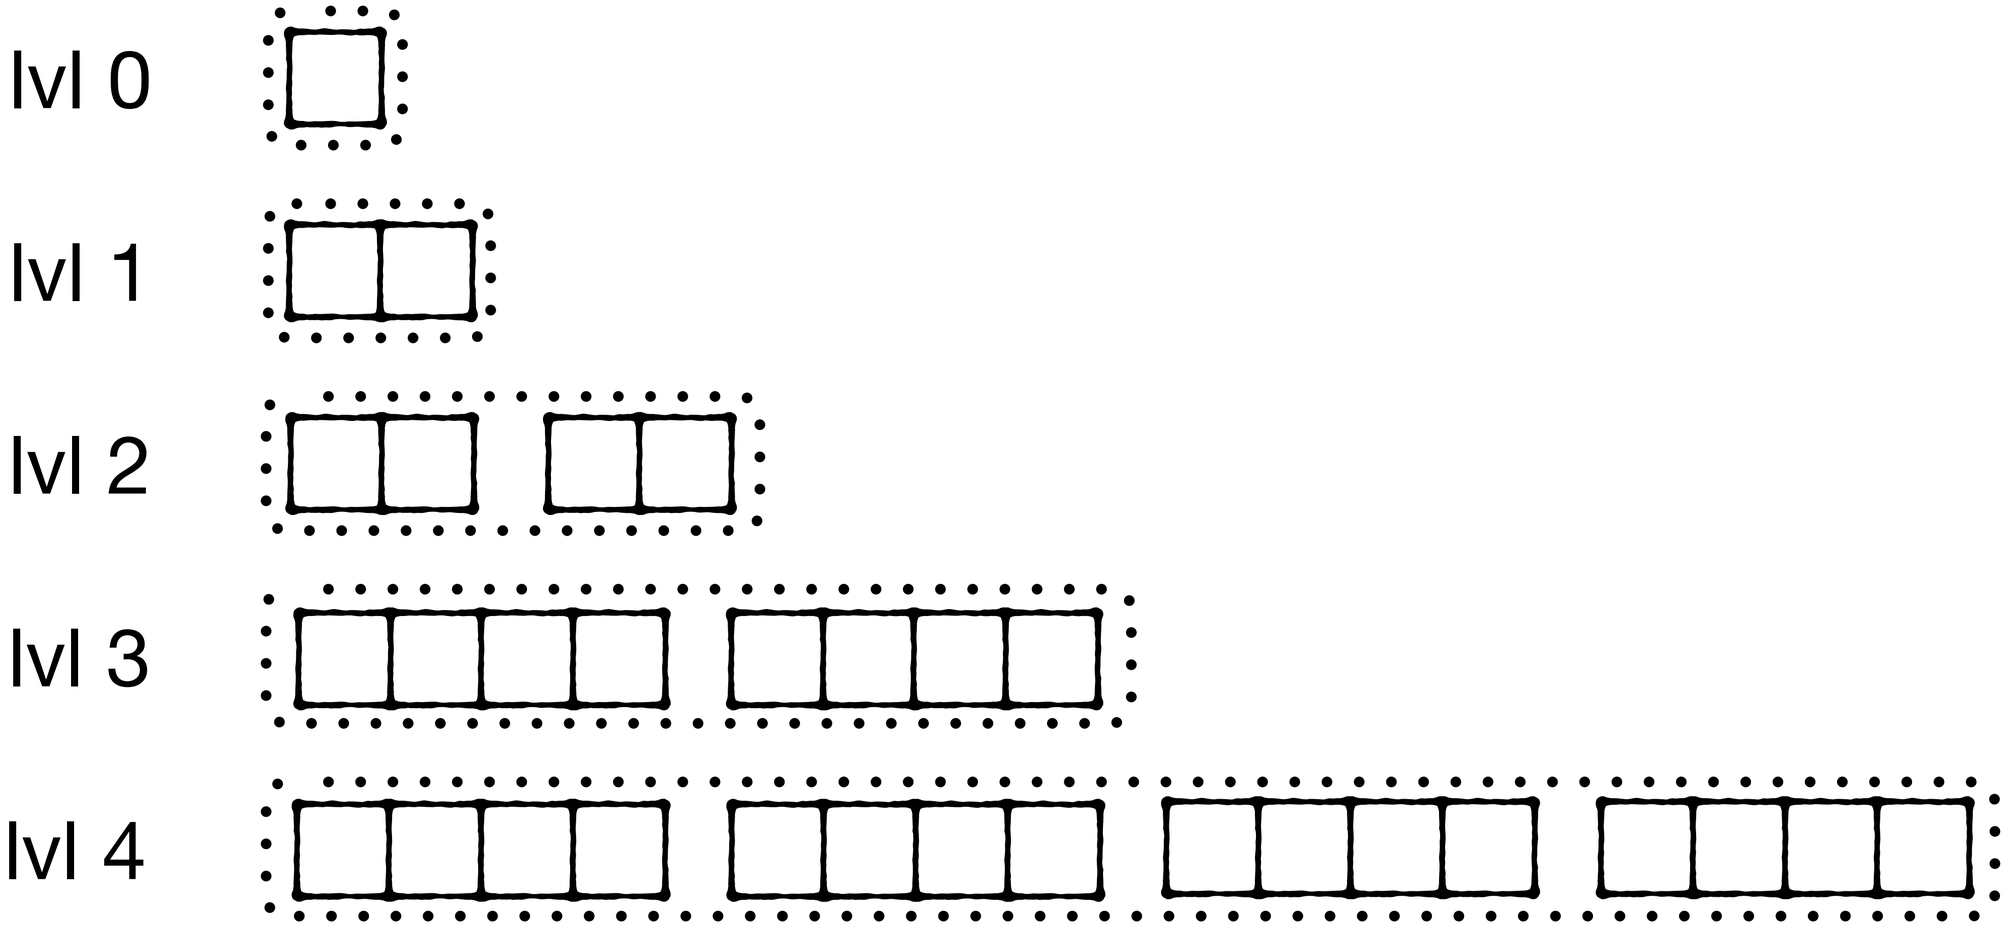

lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

levels: $\log_2 N$

 # data blocks?

lvl 0

lvl 1

lvl 2

lvl 3

lvl 4

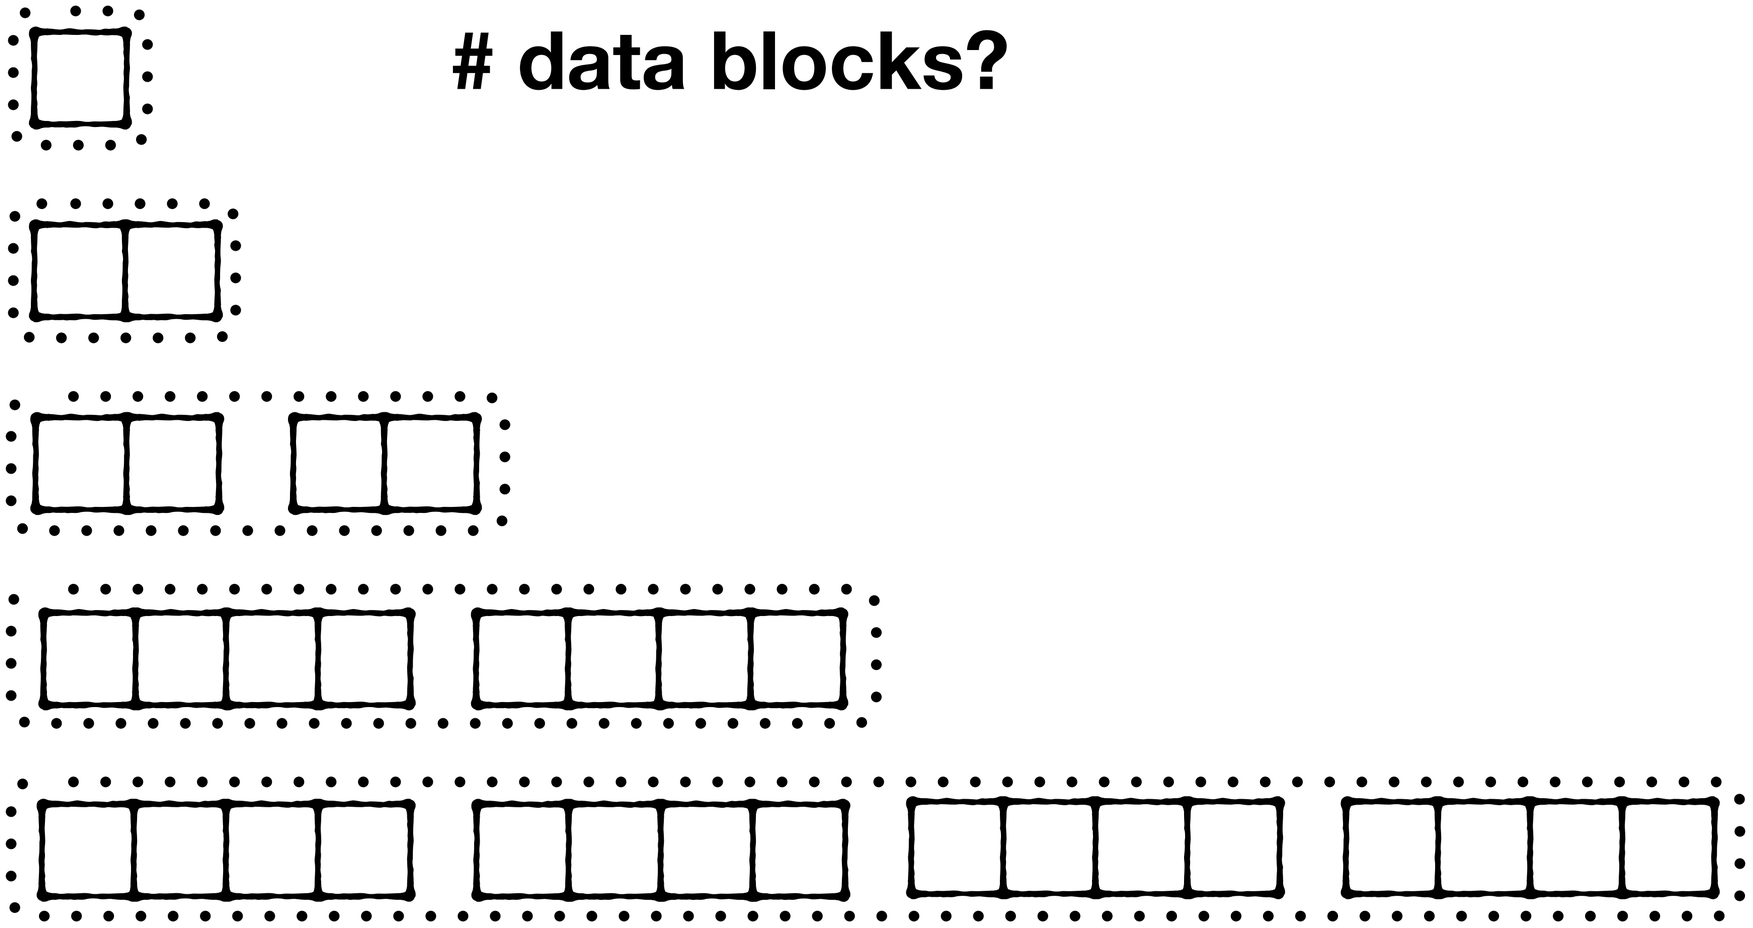

lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

levels: $\log_2 N$

 # data blocks?

Most are here

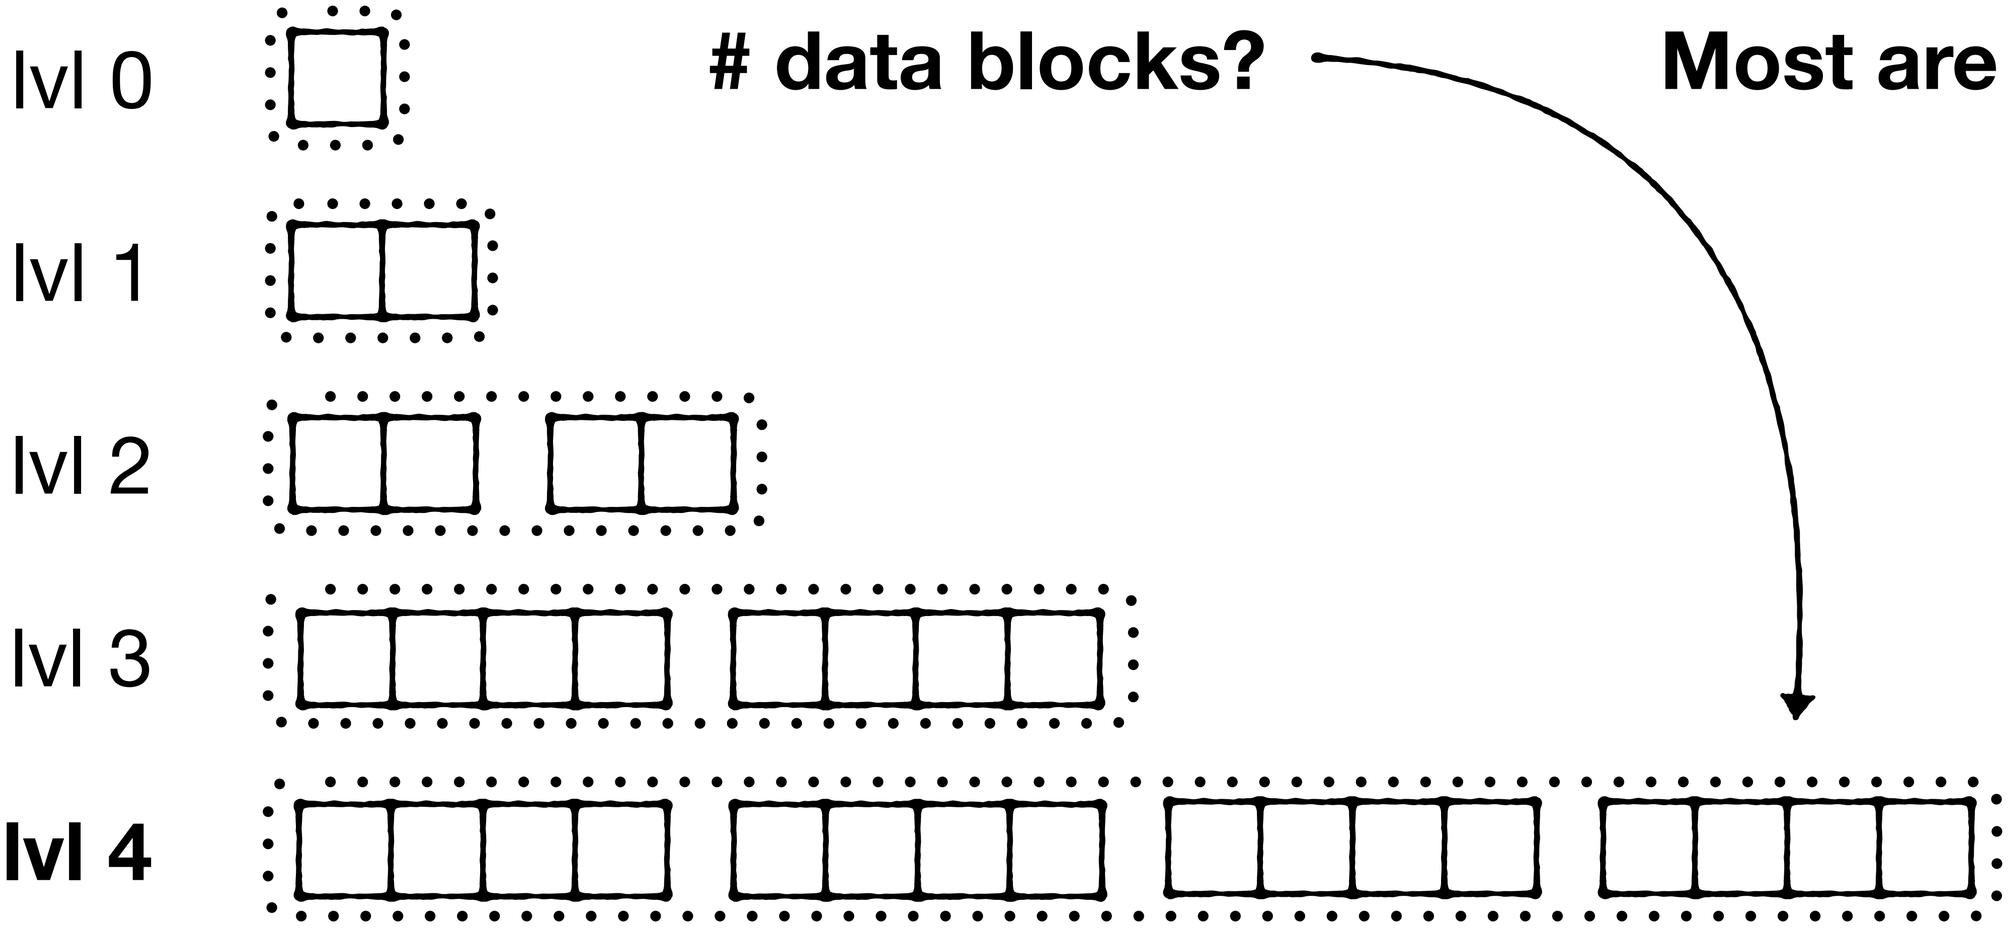

lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

levels: **$\log_2 N$**

 # data blocks? $2^{\lfloor K/2 \rfloor}$

lvl 0

lvl 1

lvl 2

lvl 3

lvl 4

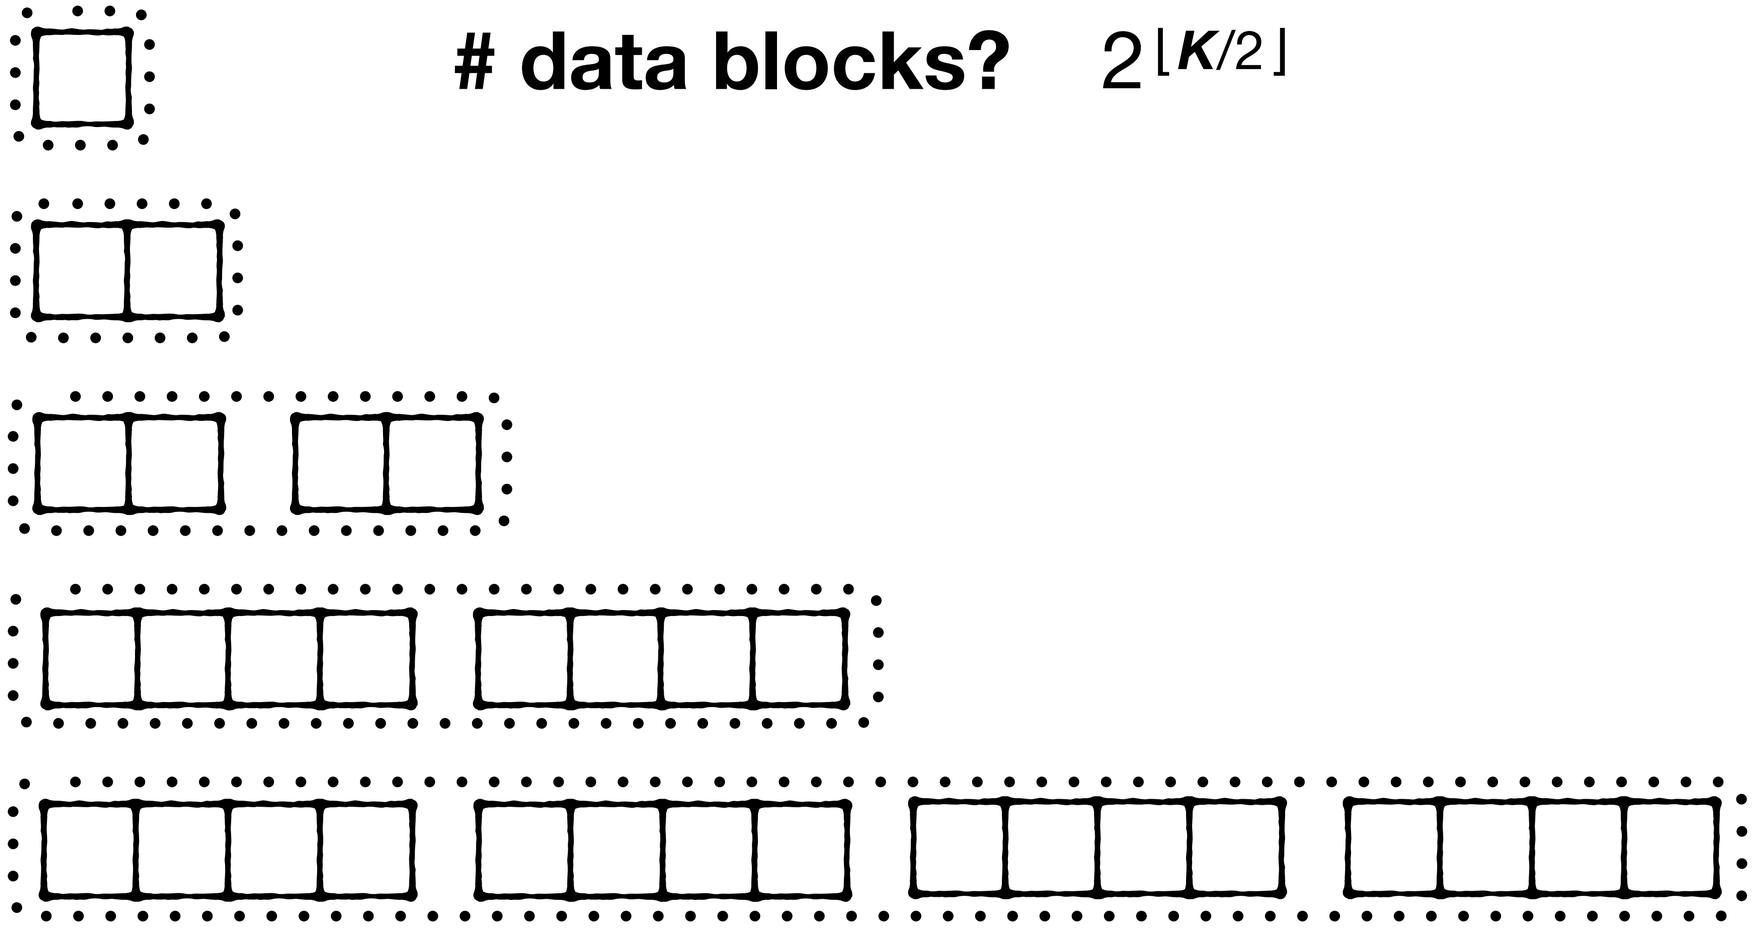

lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

levels: **$\log_2 N$**

 # data blocks? $2^{\lfloor \log N/2 \rfloor}$

lvl 0

lvl 1

lvl 2

lvl 3

lvl 4

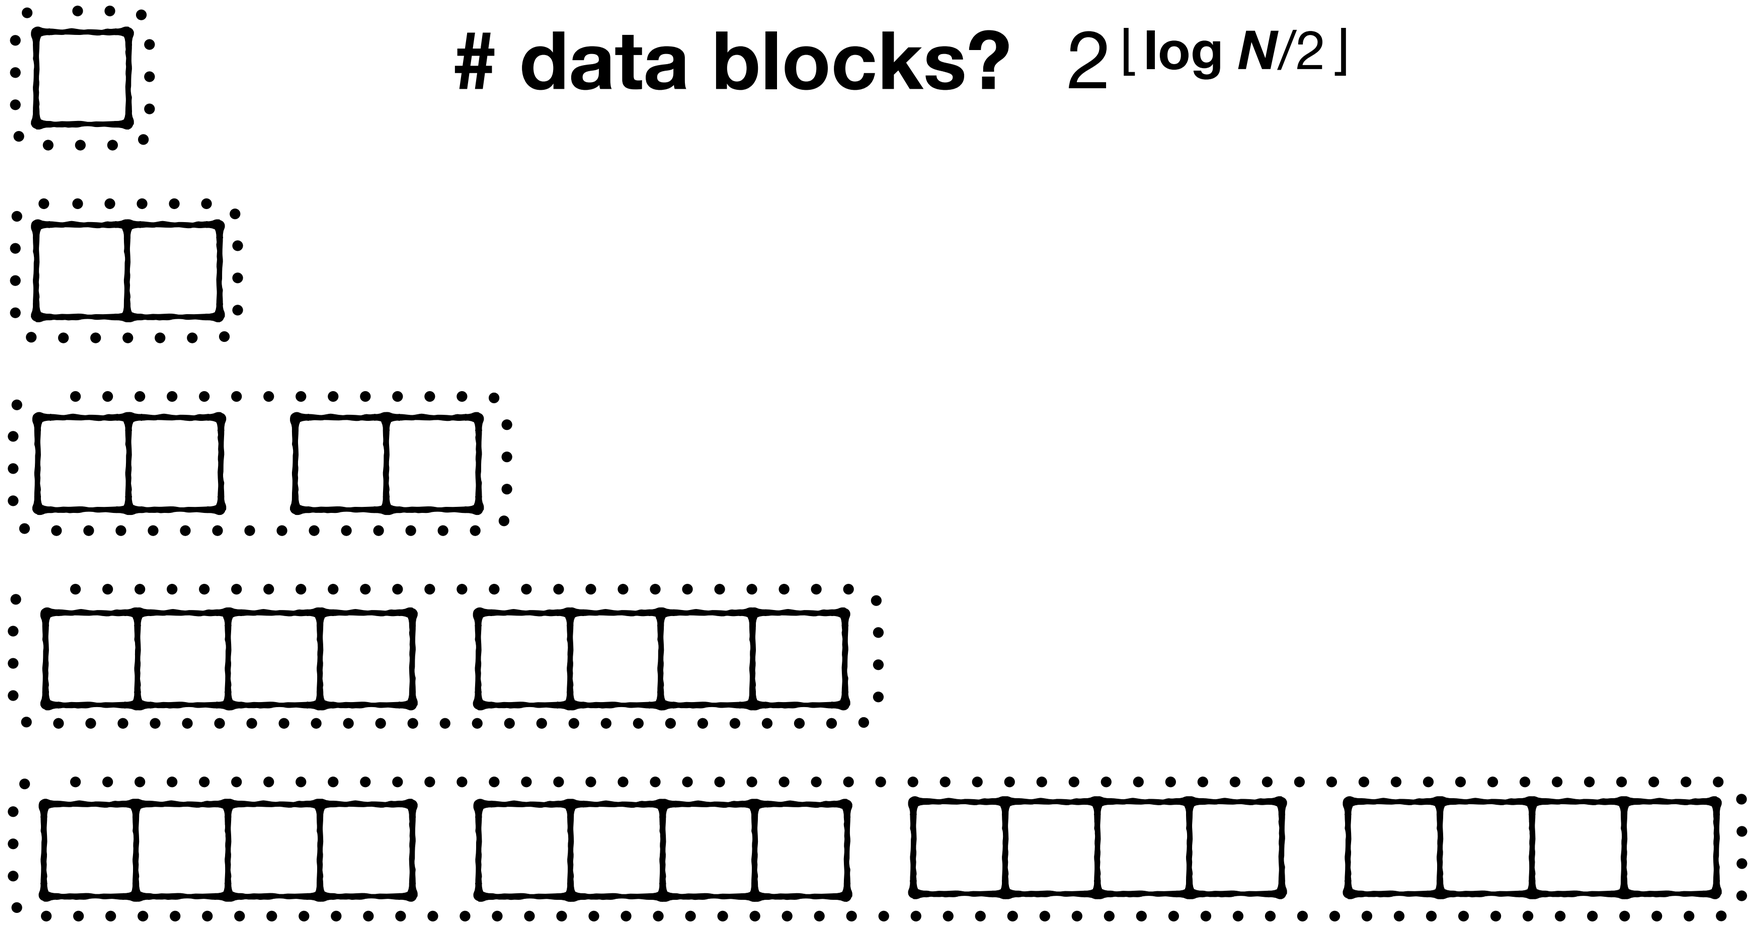

lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

levels: $\log_2 N$

 # data blocks? O(√N)

lvl 0

lvl 1

lvl 2

lvl 3

lvl 4

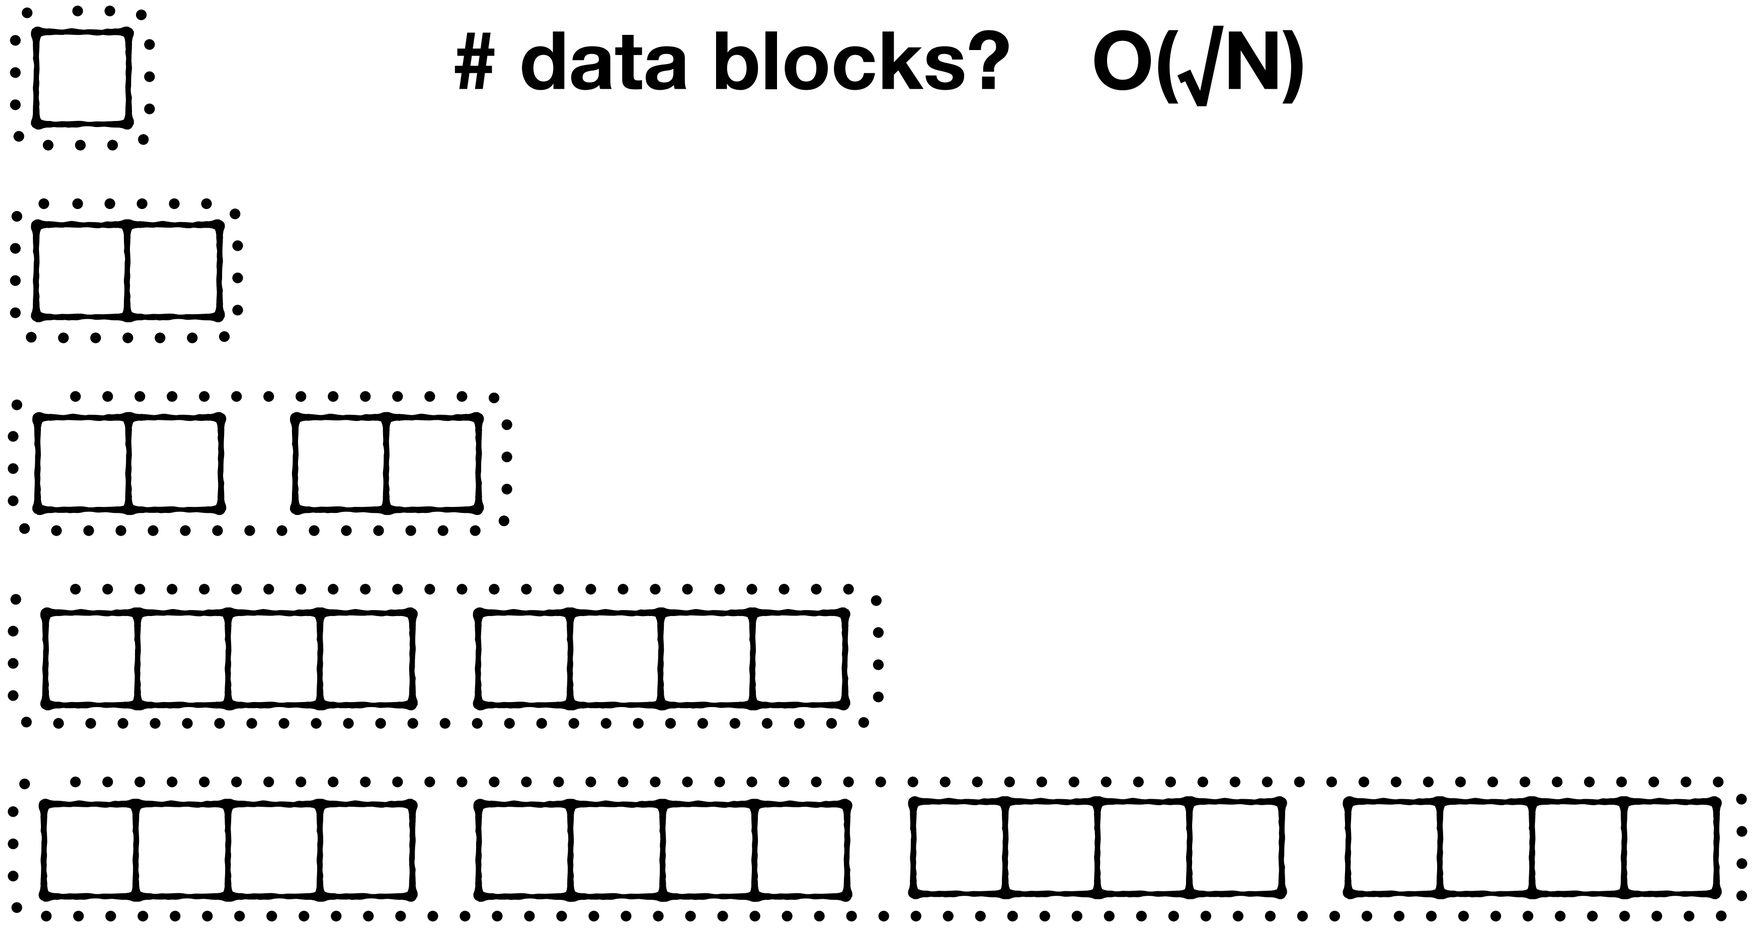

lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

levels: $\log_2 N$

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

 # slots in largest block?

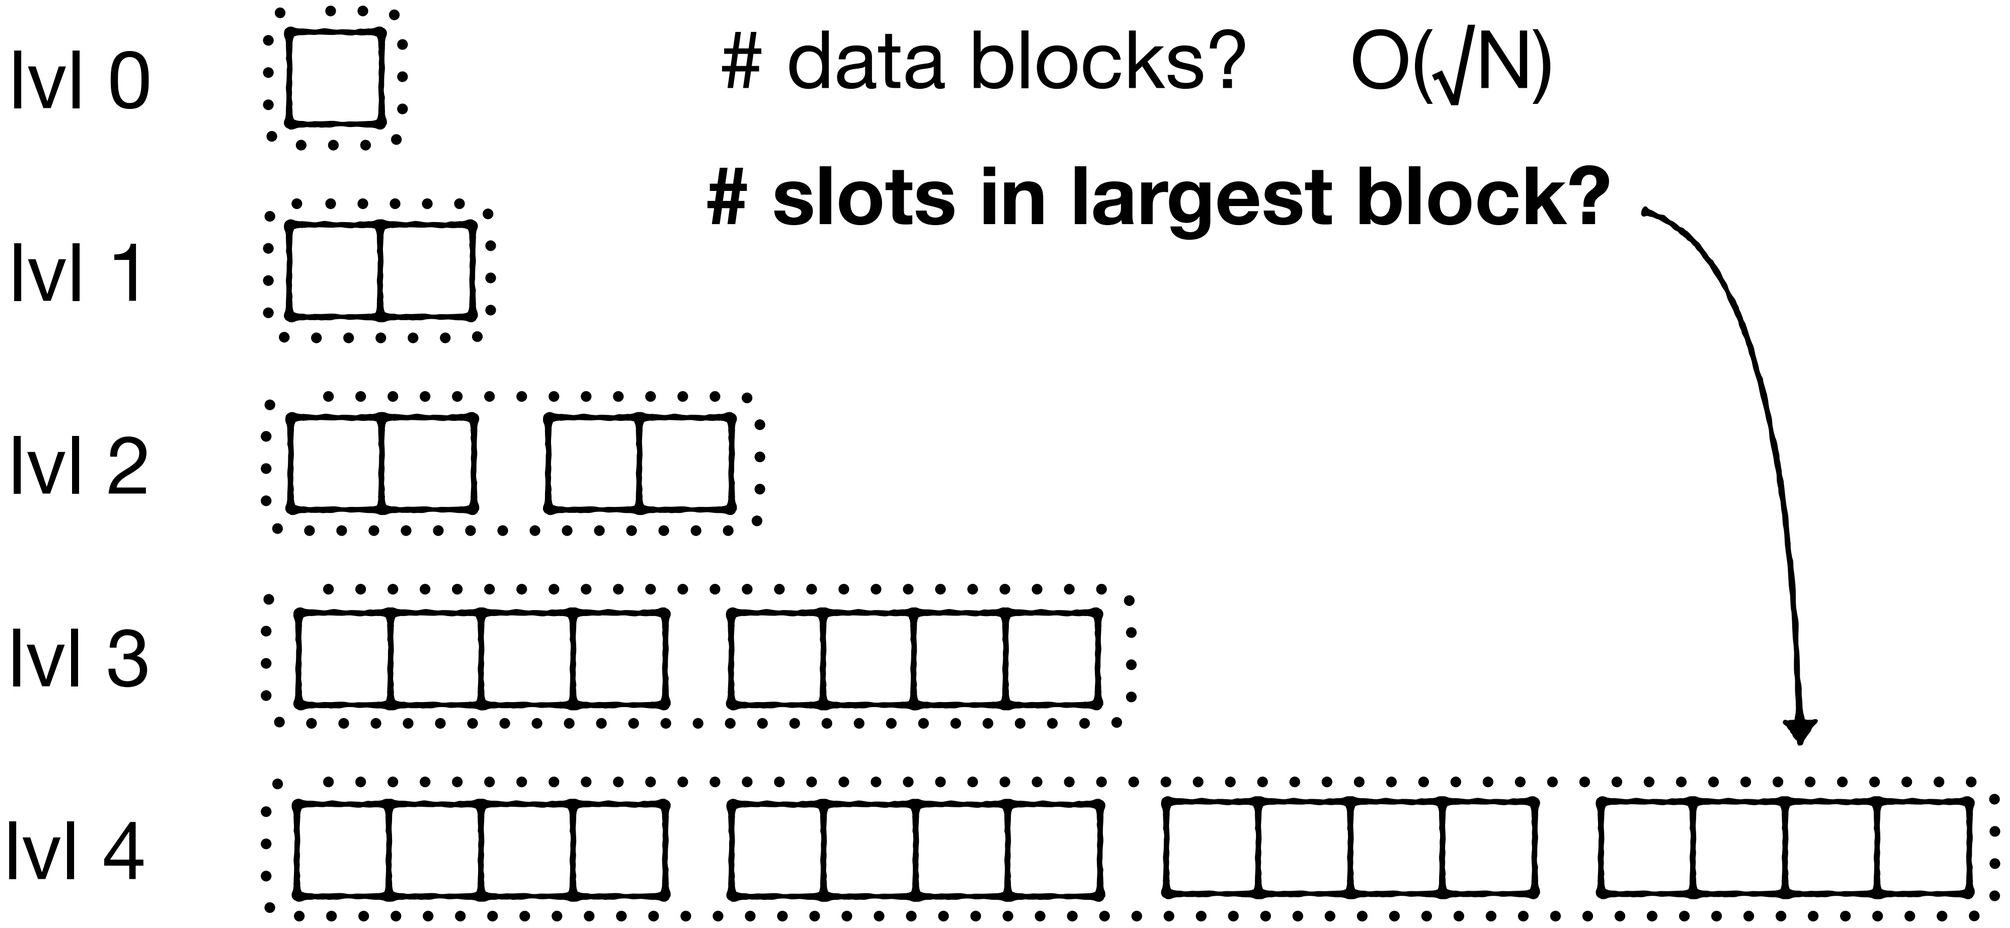

lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

levels: **$\log_2 N$**

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

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

lvl 0

lvl 1

lvl 2

lvl 3

lvl 4

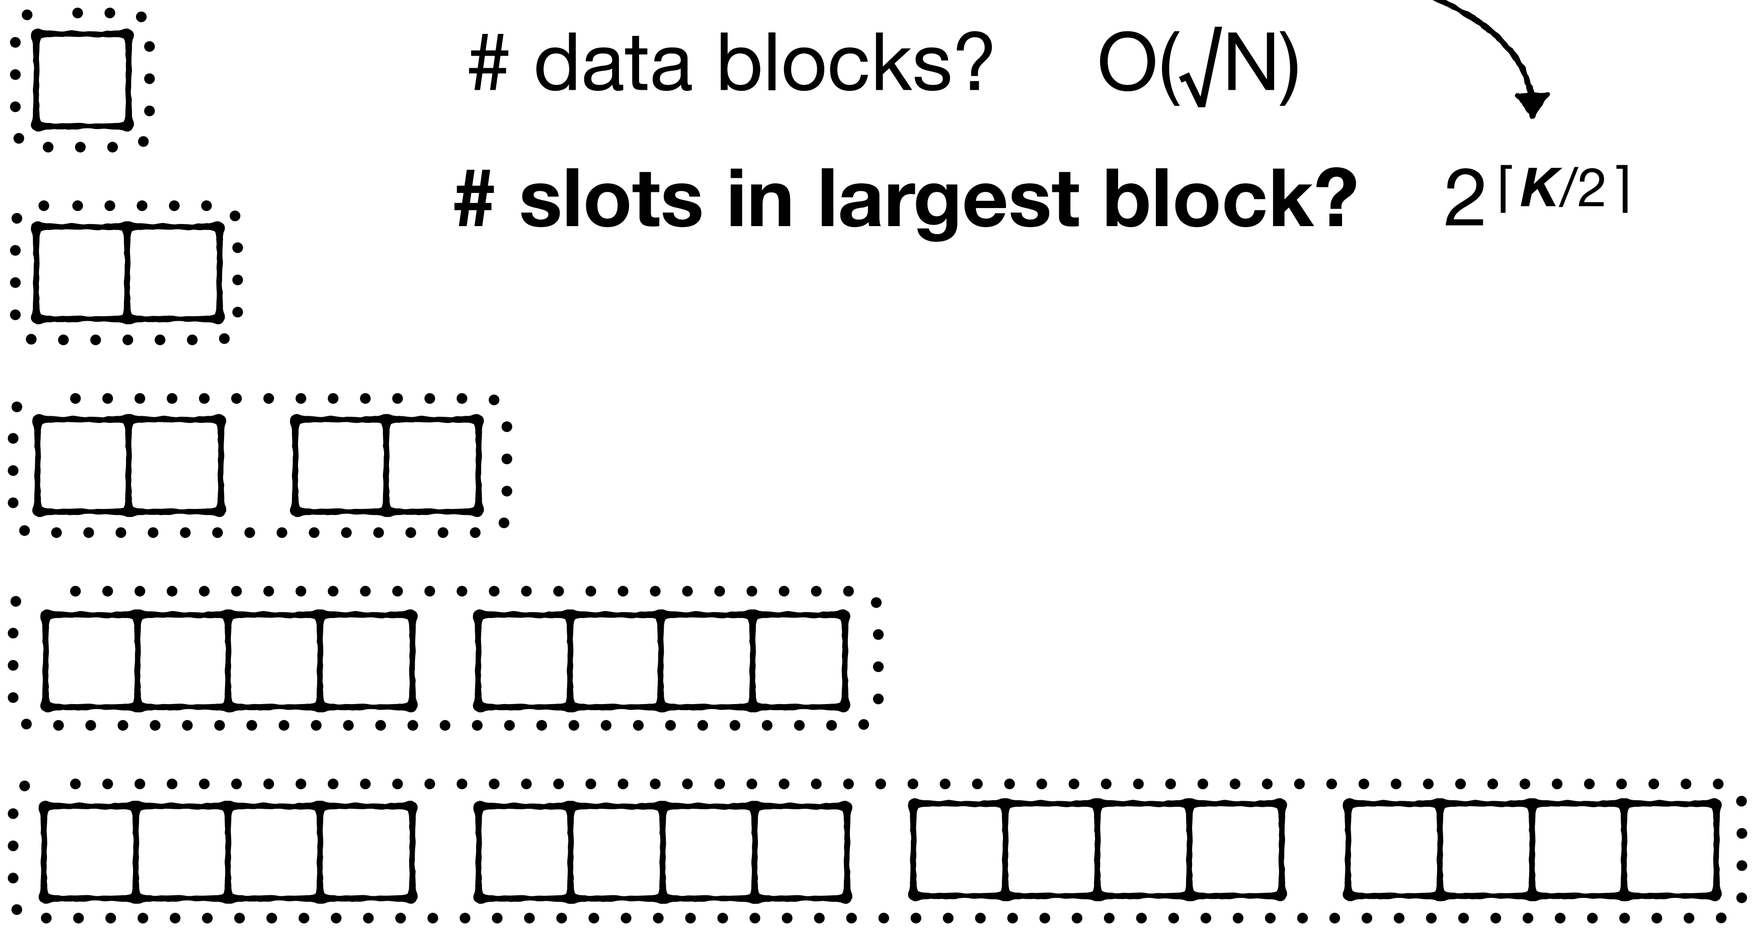

lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

levels: $\log_2 N$

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

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

lvl 0

lvl 1

lvl 2

lvl 3

lvl 4

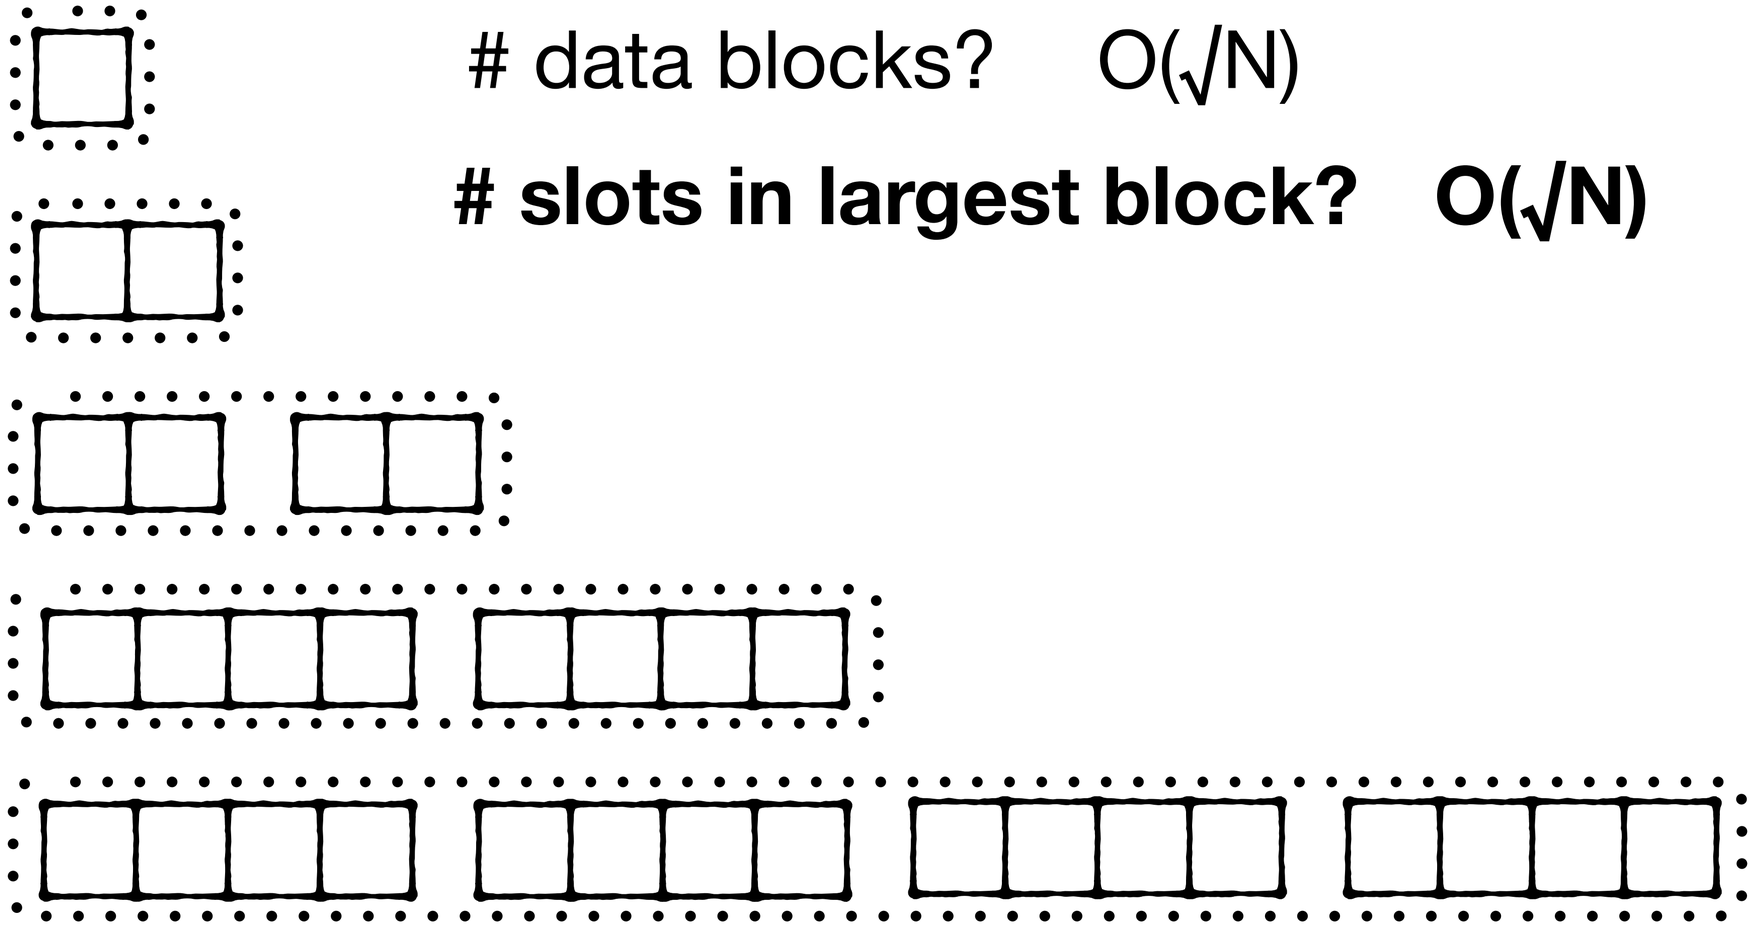

lvl i contains $2^{\lfloor K/2 \rfloor}$ blocks, each with $2^{\lceil K/2 \rceil}$ slots

levels: $\log_2 N$

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

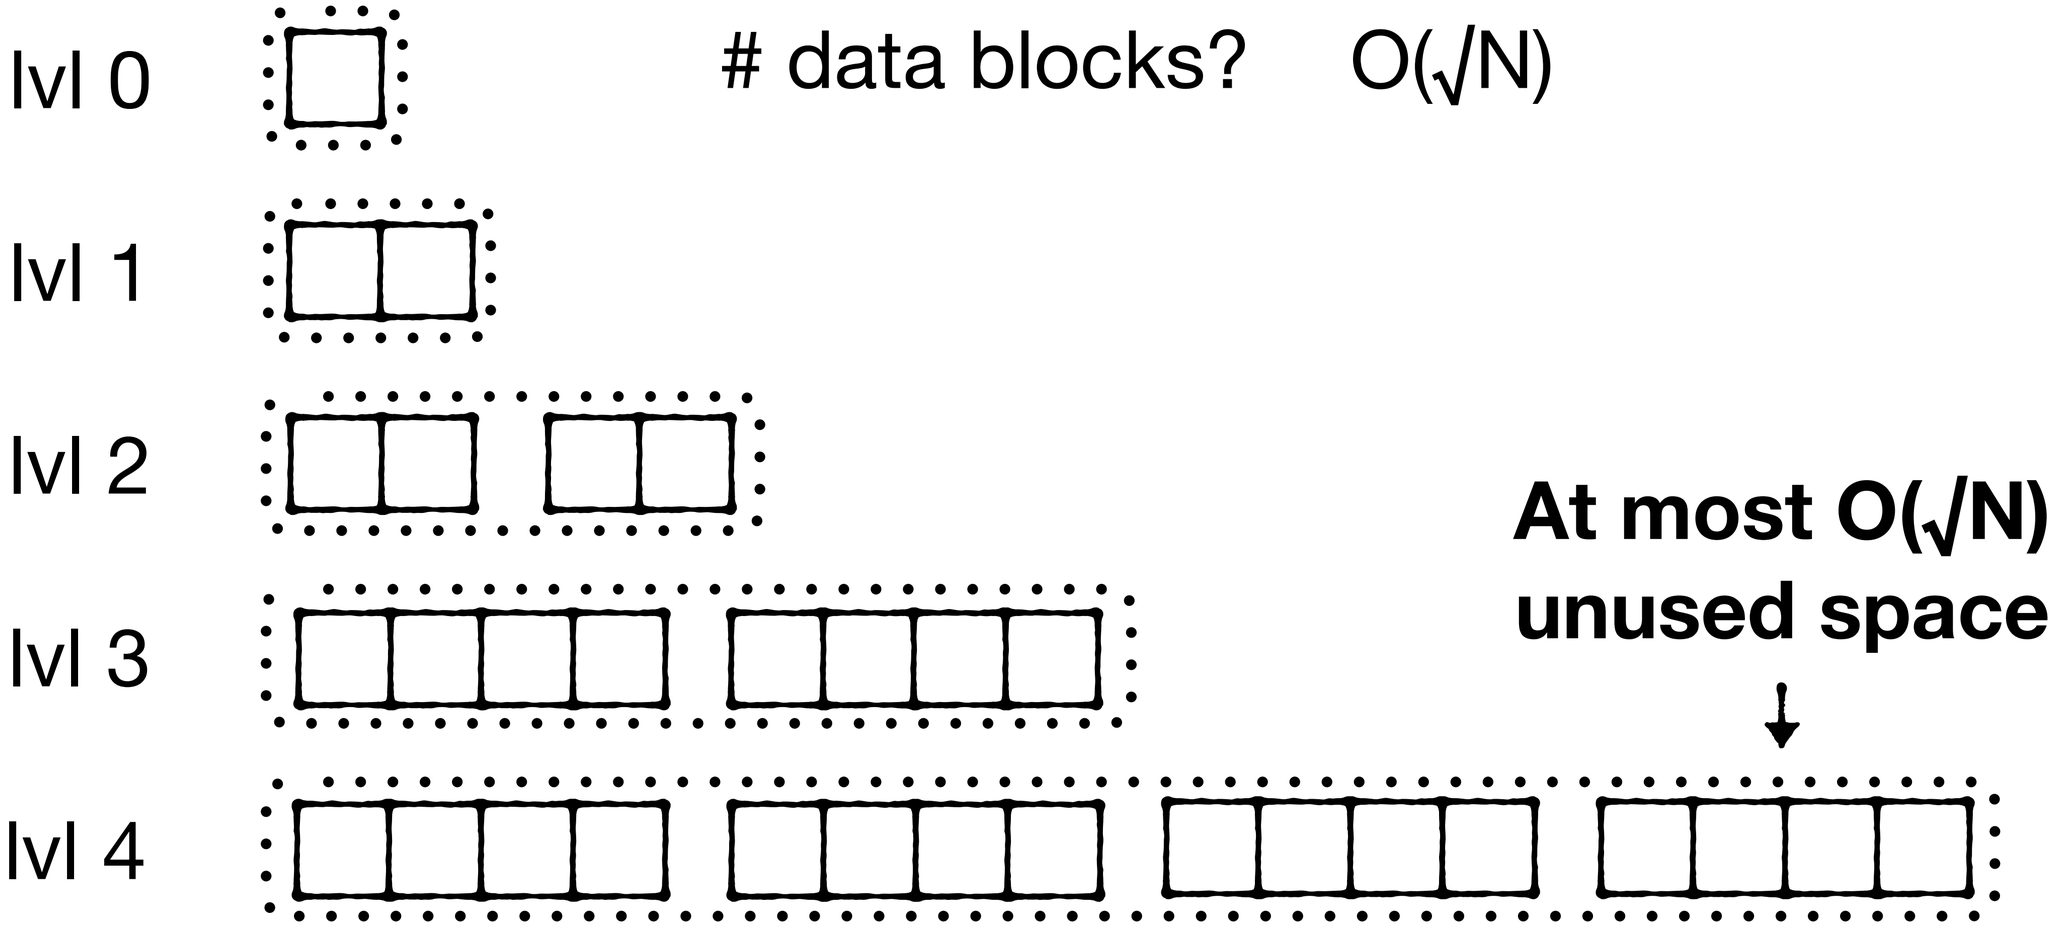

Directory with O(√N) pointers

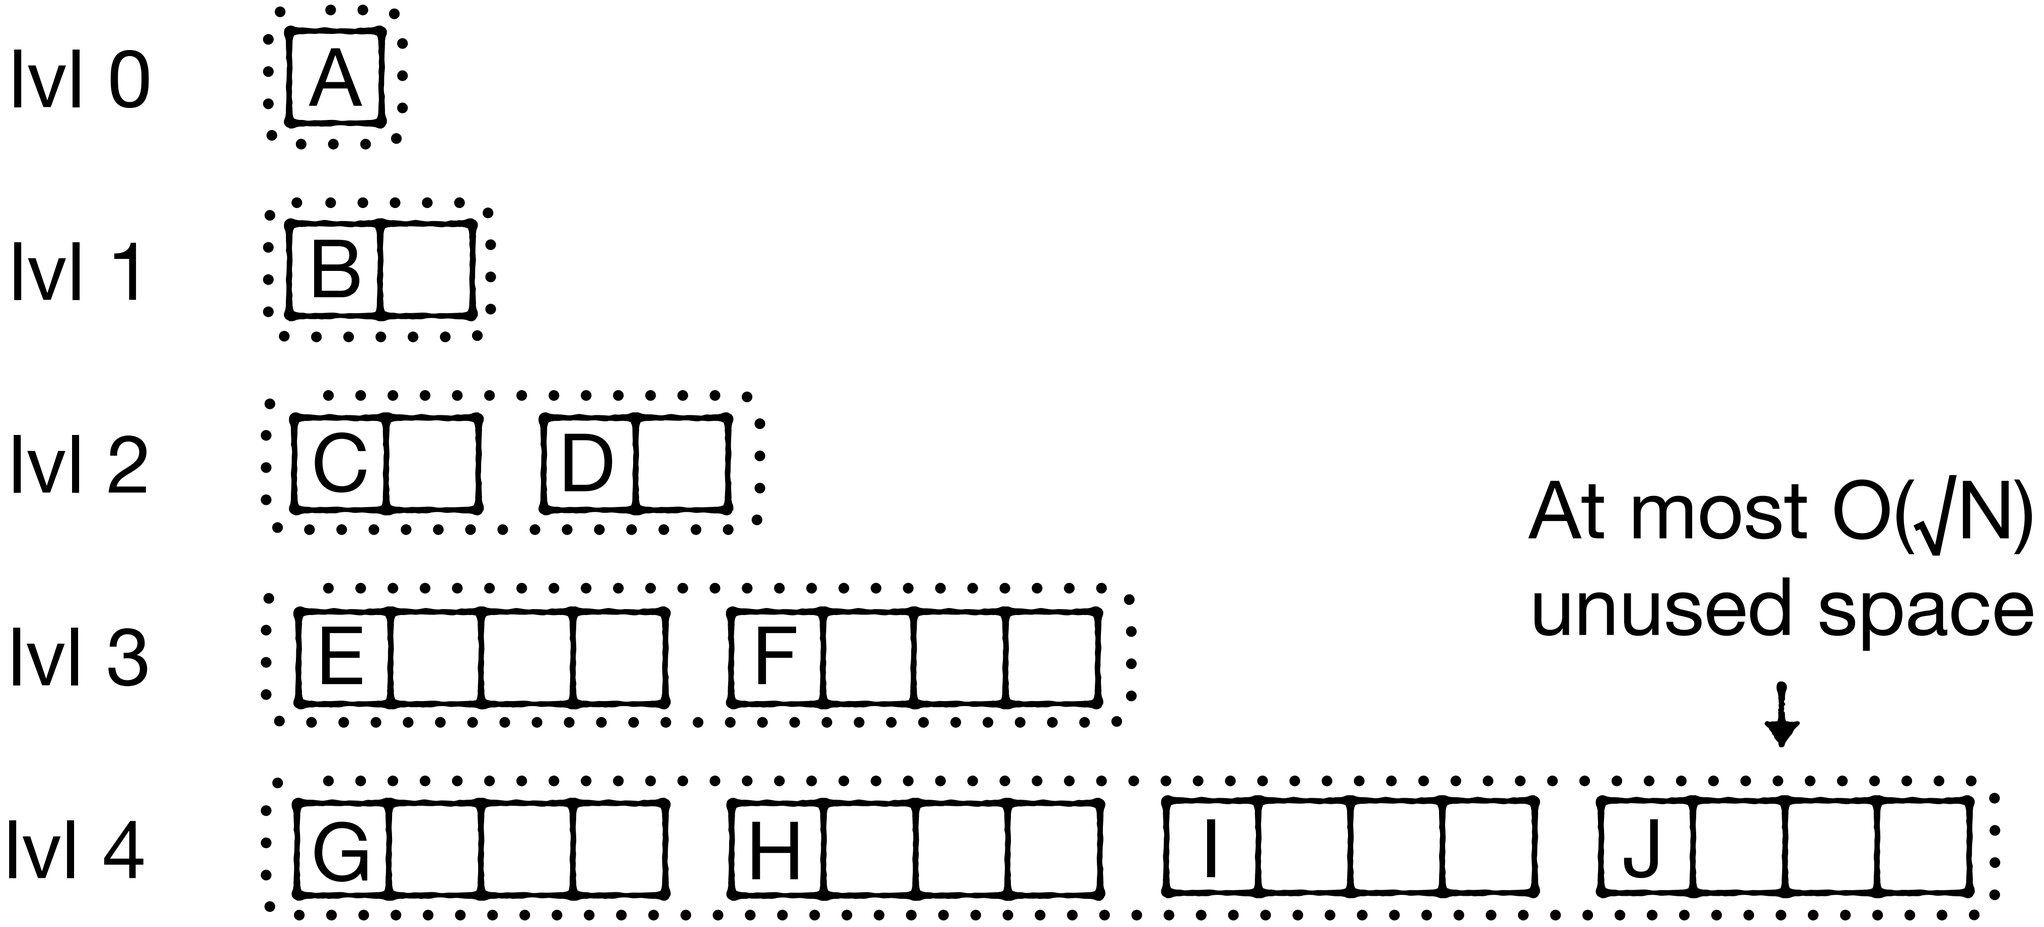

At most half O(√N) unused space

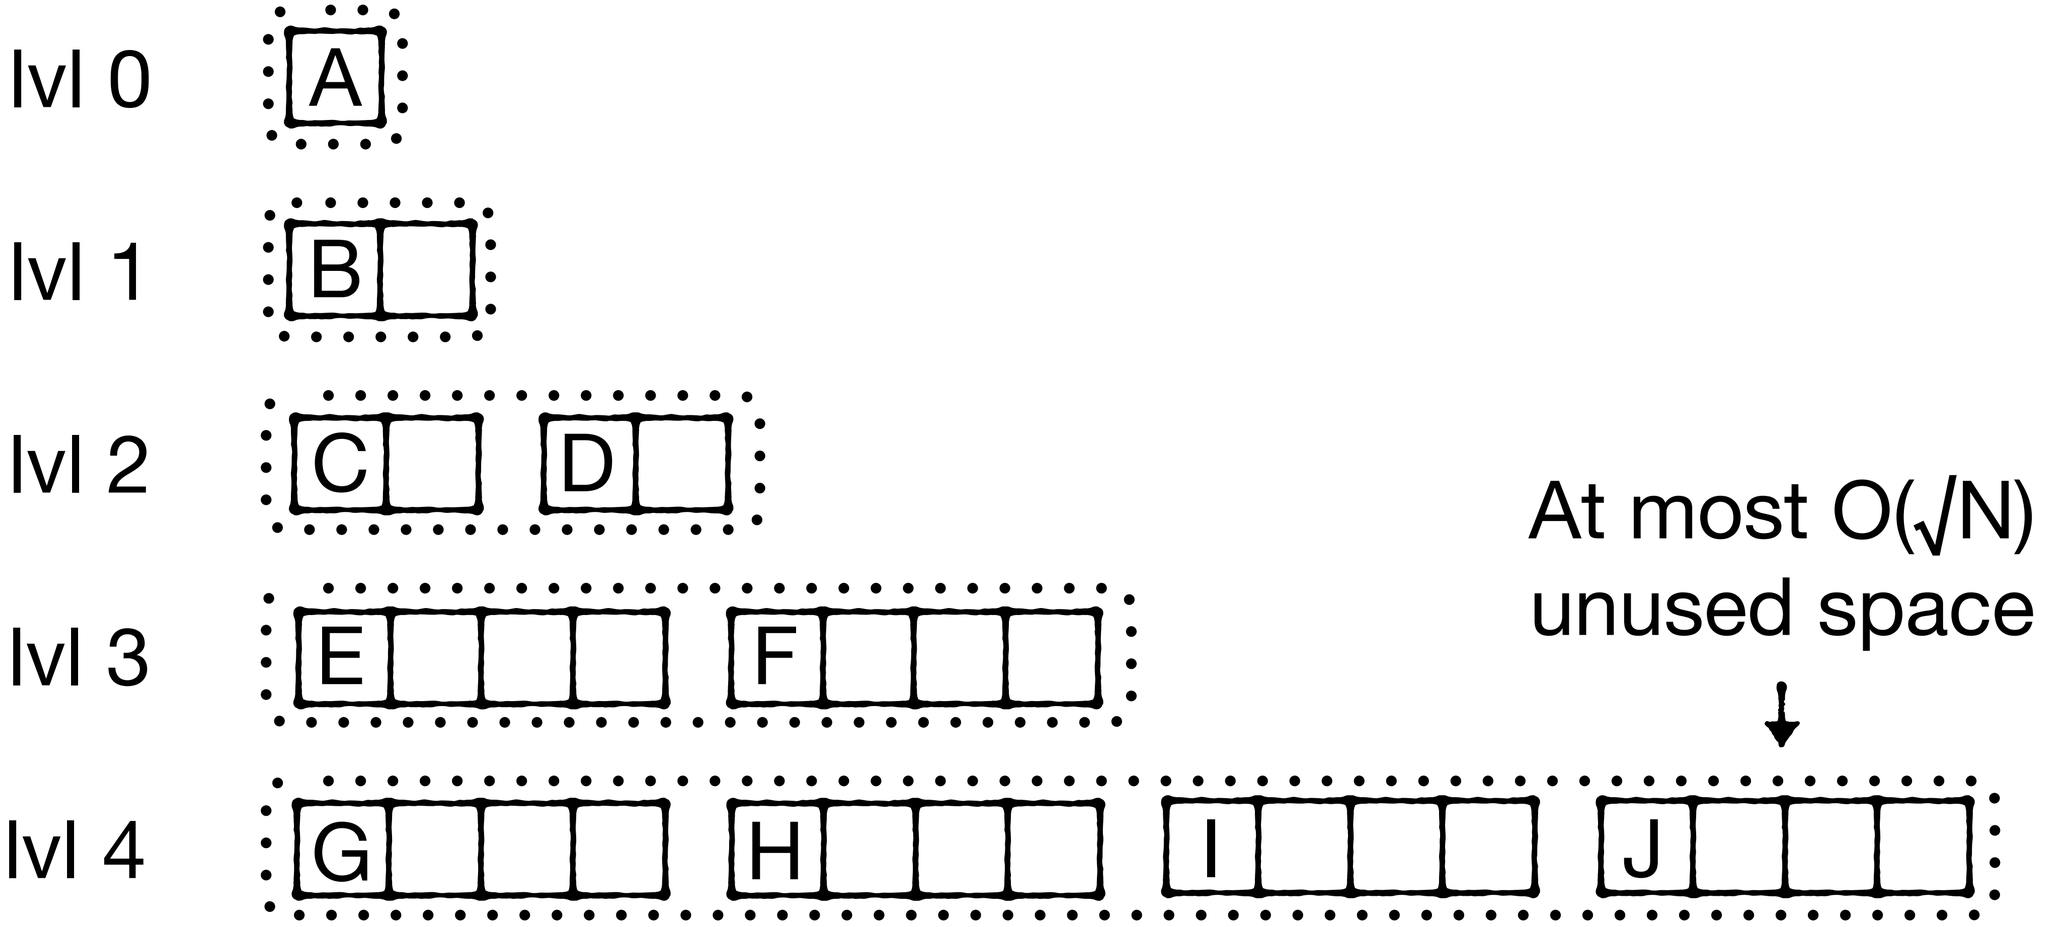

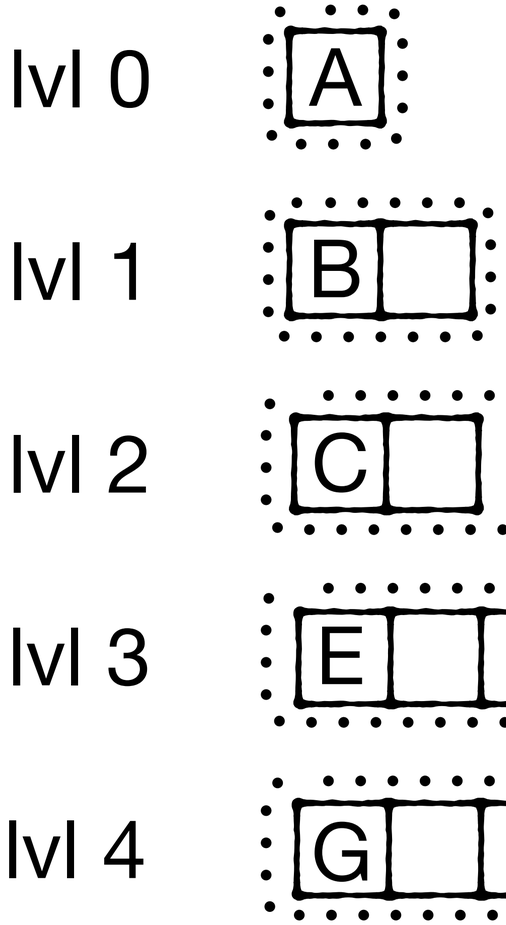

Max extra space: $O(\sqrt{N}) + O(\sqrt{N})$ **= $O(\sqrt{N})$**

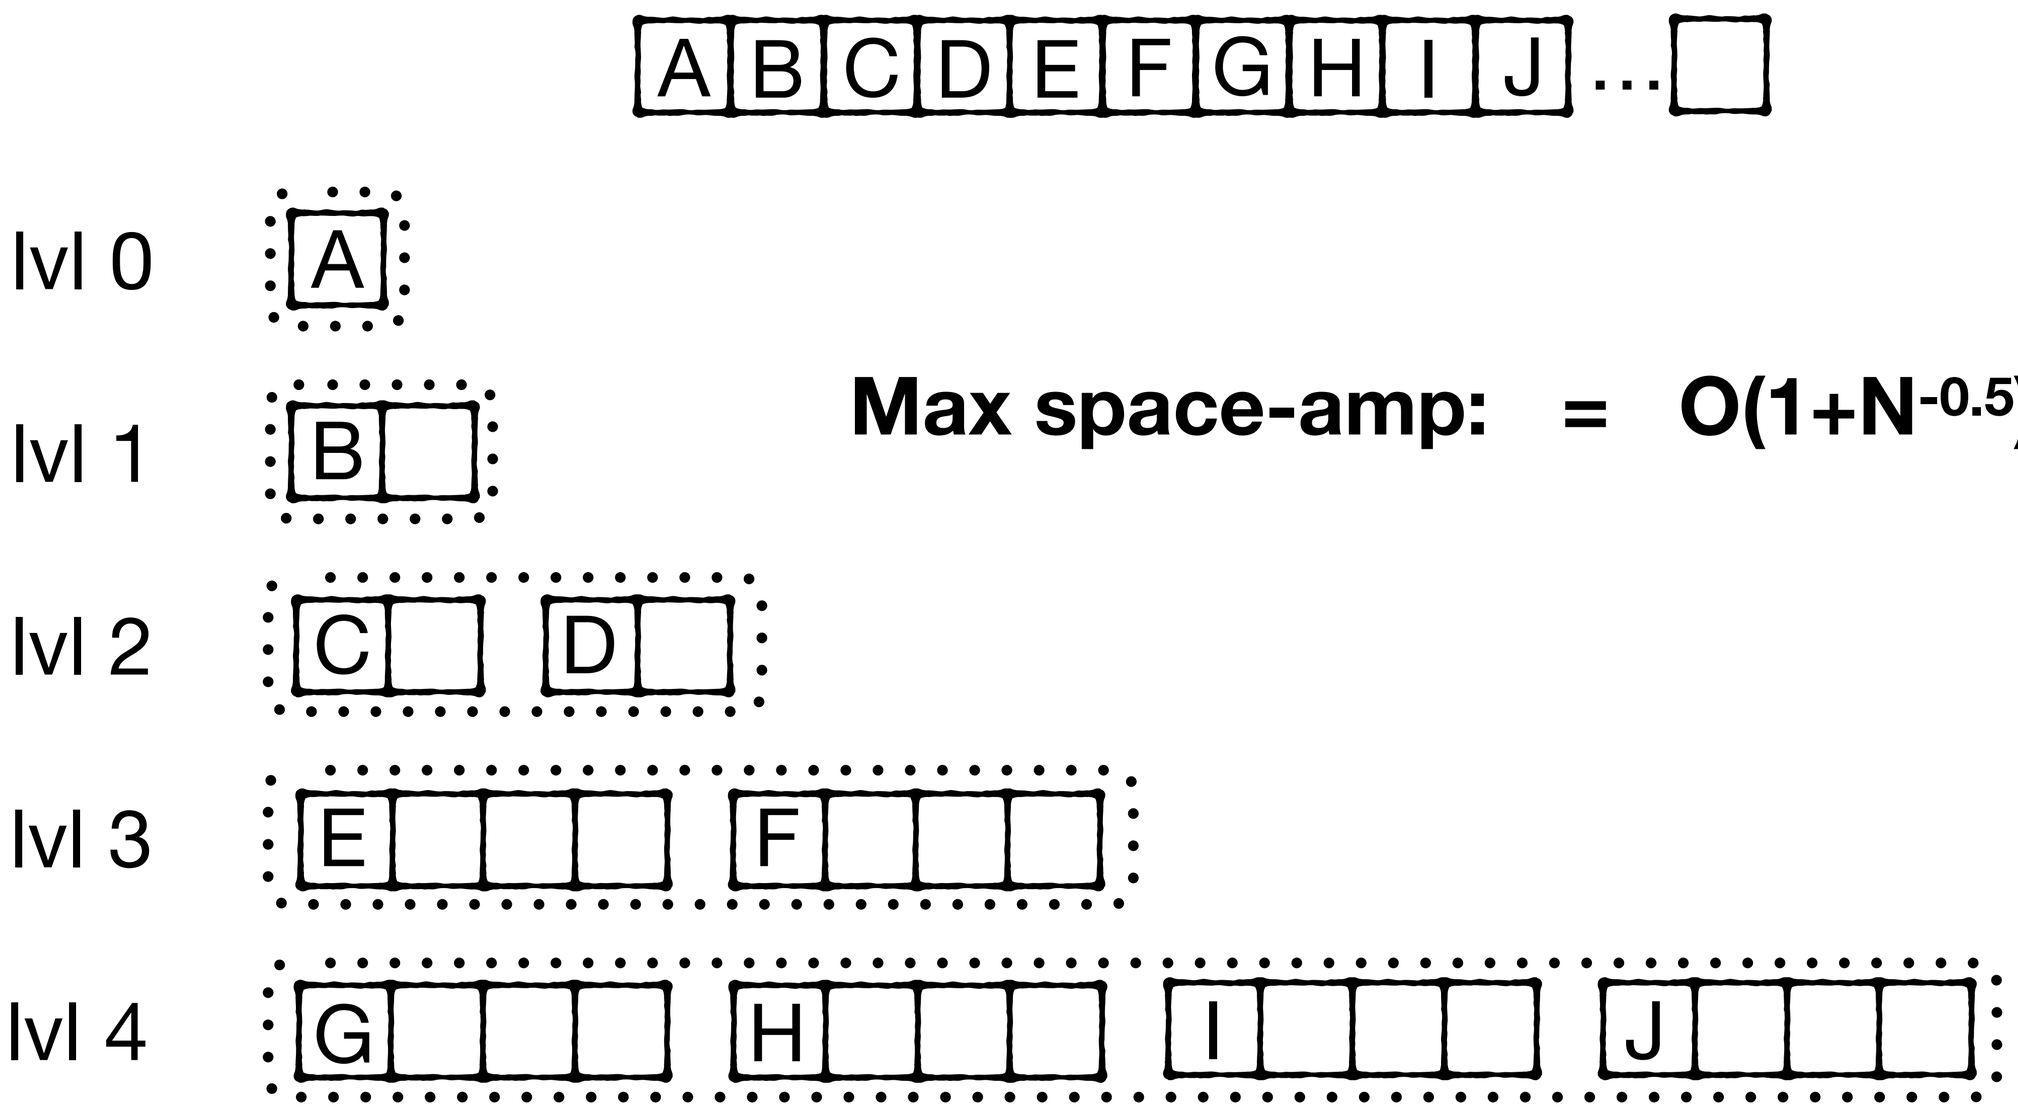

Max space-amp: = $O(1+N^{-0.5})$

How to access slot in O(1) time?

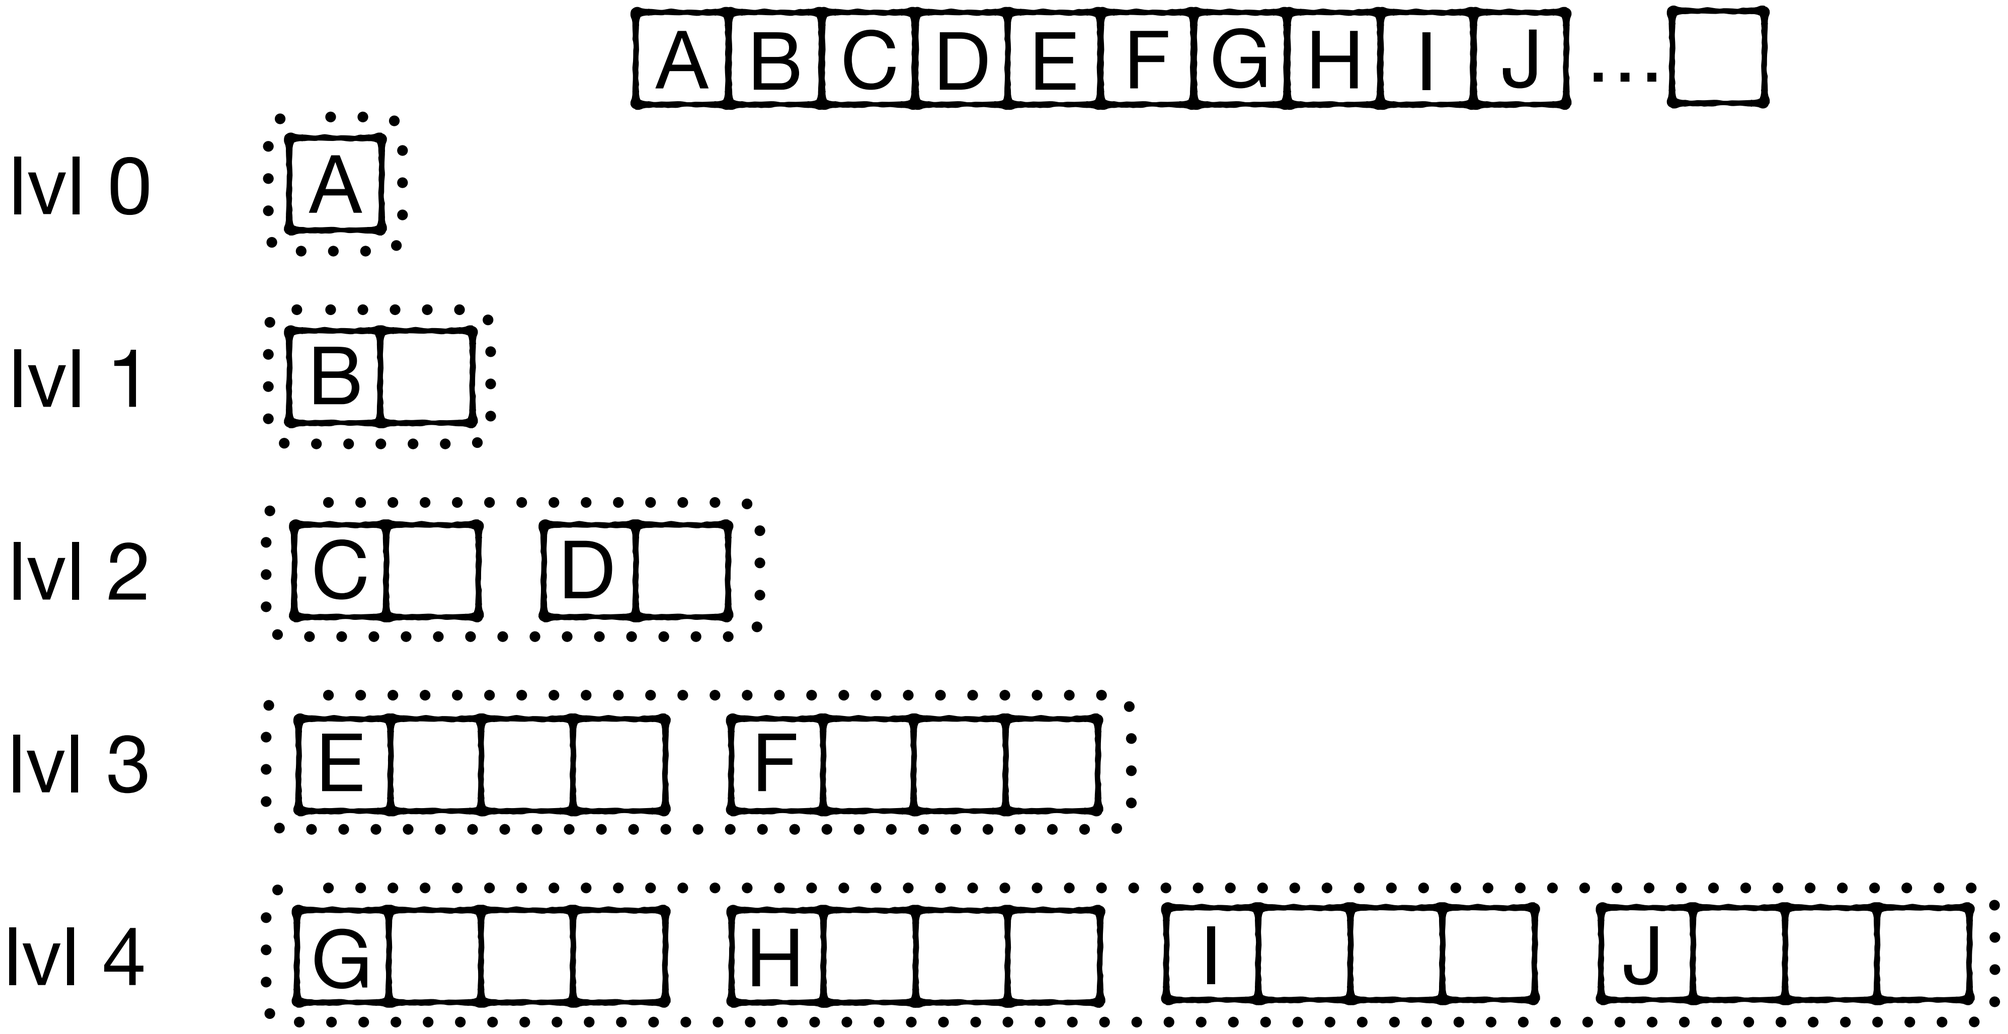

get(12)

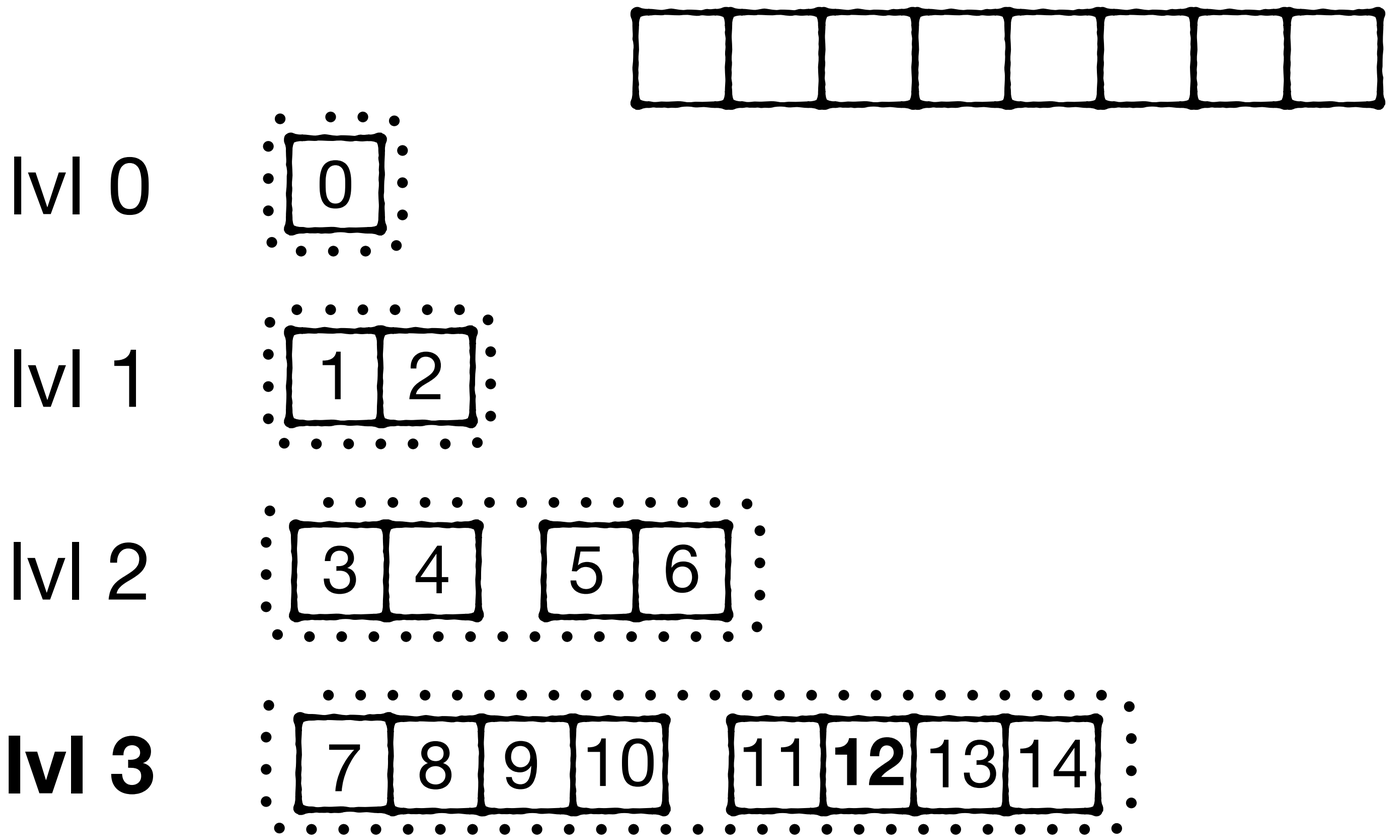

get(12)

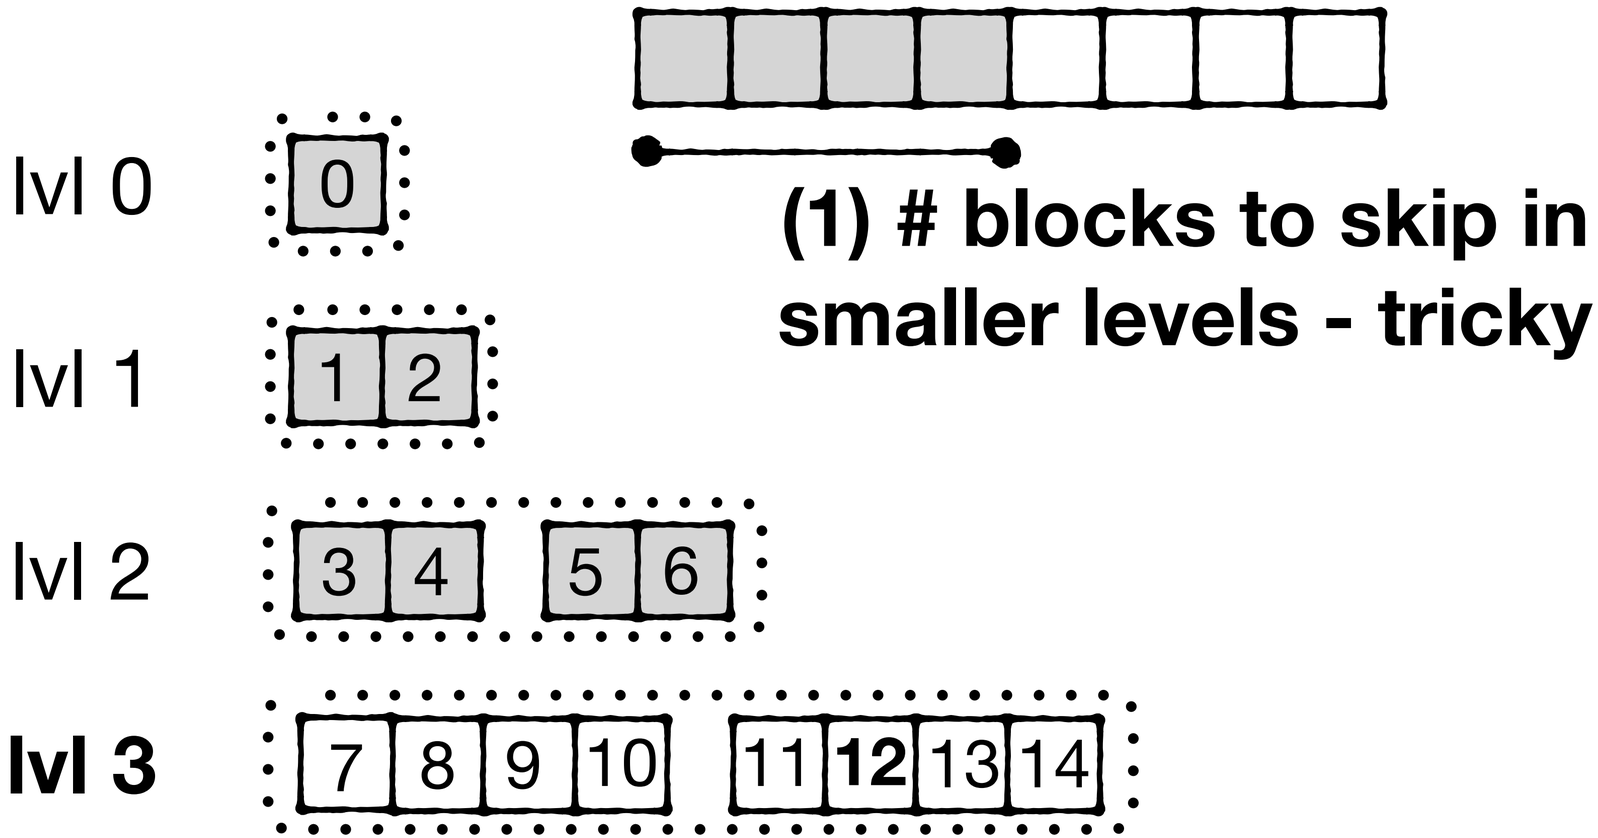
lvl 0
0
lvl 1
1
2
lvl 2
3
4
5
6
lvl 3
7
8
9
10
11
12
13
14
(1) # blocks to skip in smaller levels - tricky

get(12)

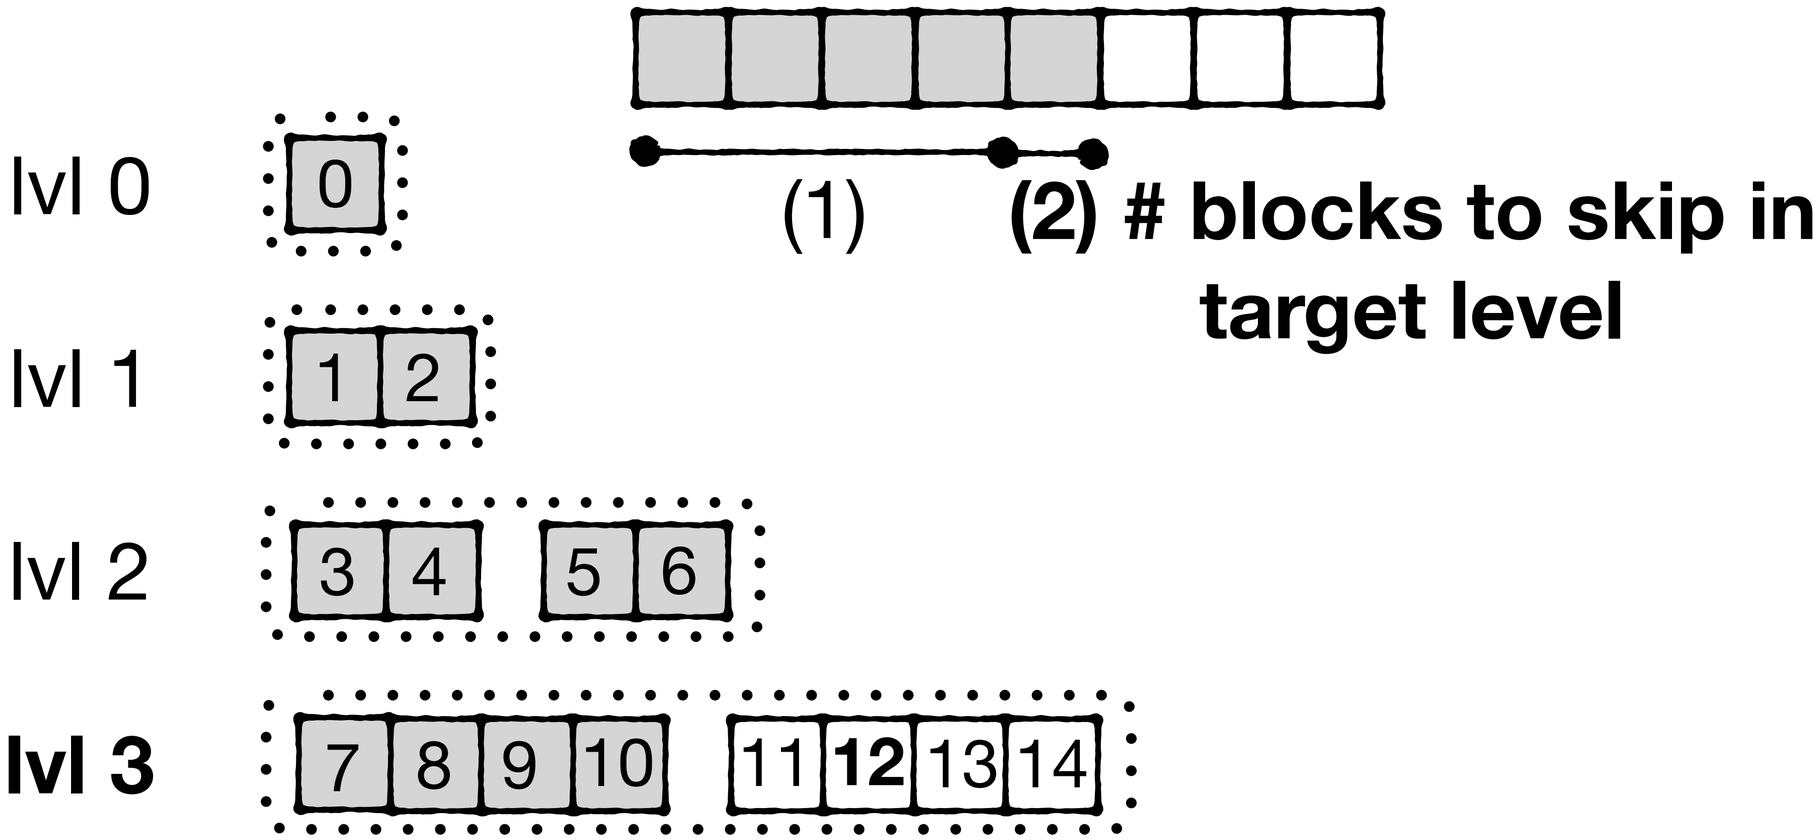
lvl 0
0
lvl 1
1
2
lvl 2
3
4
5
6
lvl 3
7
8
9
10
11
12
13
14
(1)
(2) # blocks to skip in target level

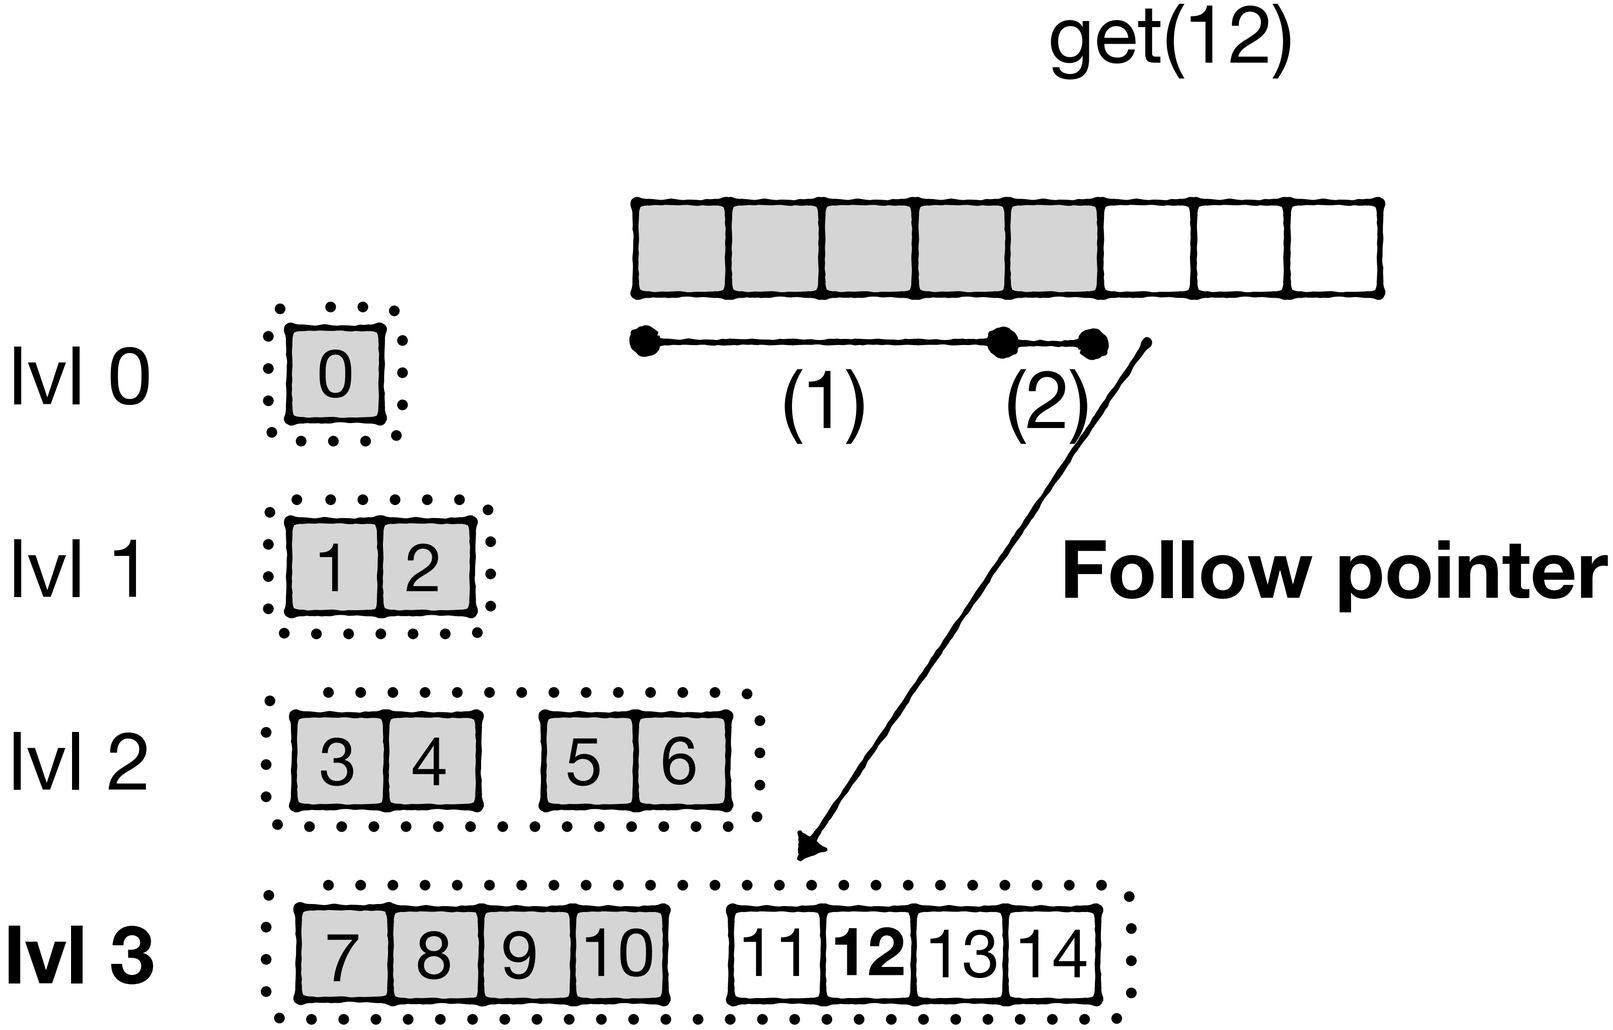
get(12)
lvl 0
0
(1)
(2)
lvl 1
1
2
Follow pointer
lvl 2
3
4
5
6
lvl 3
7
8
9
10
11
12
13
14

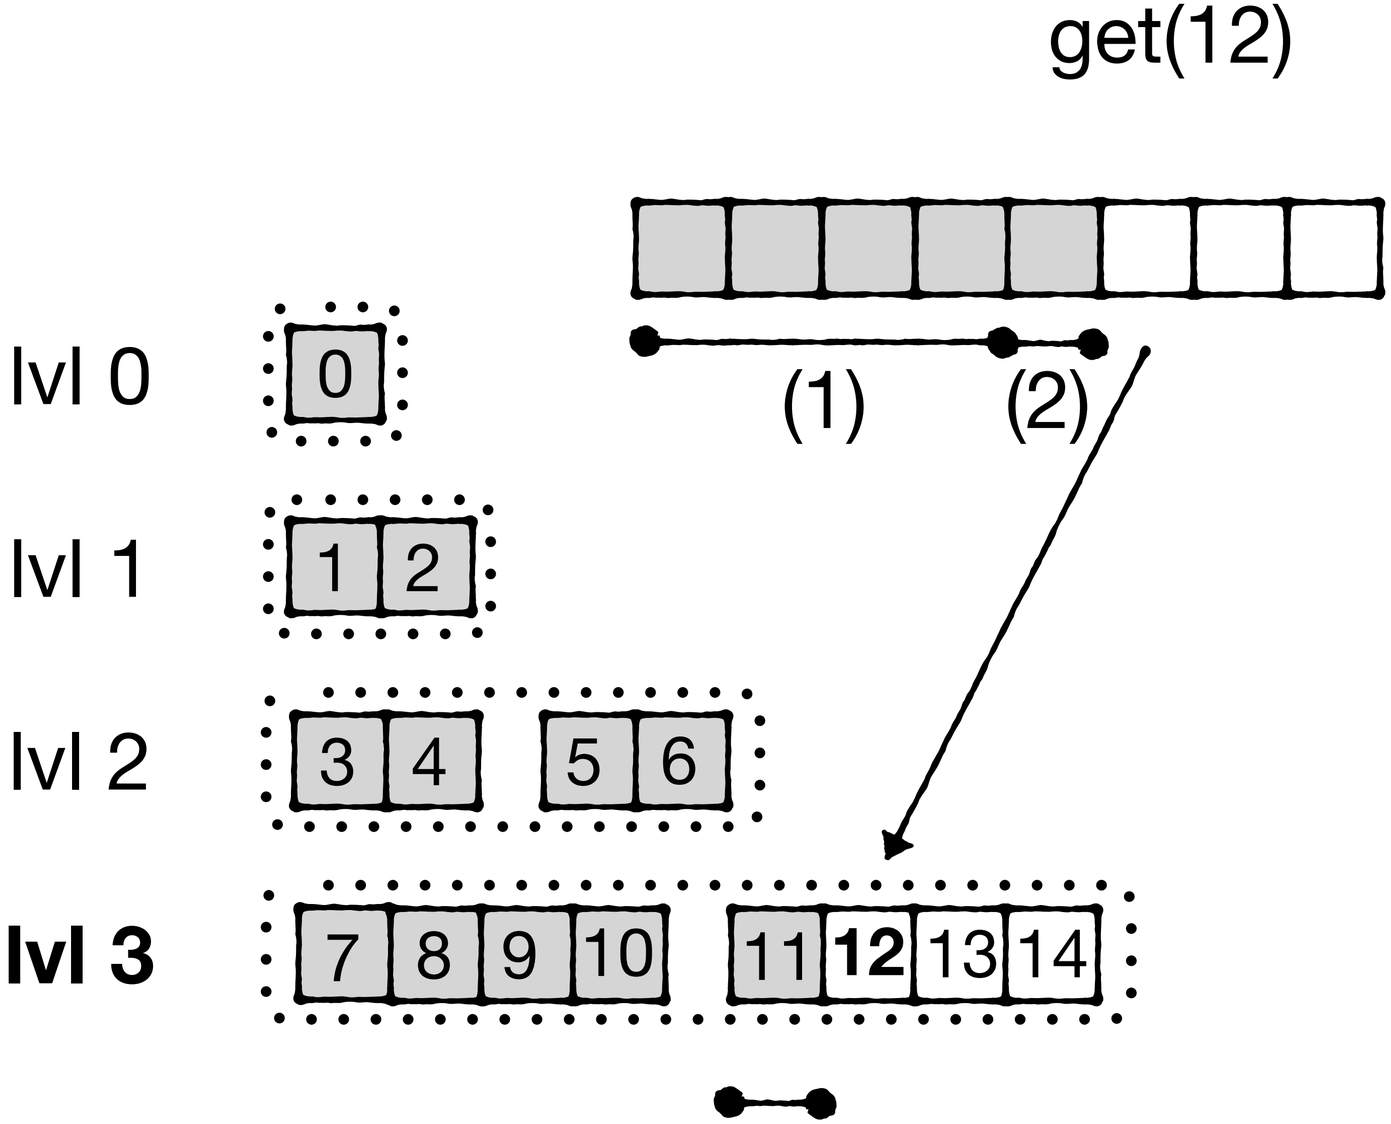
get(12)
lvl 0
0
lvl 1
1
2
lvl 2
3
4
5
6
lvl 3
7
8
9
10
11
12
13
14
(1)
(2)
(3) # slots to skip within target block

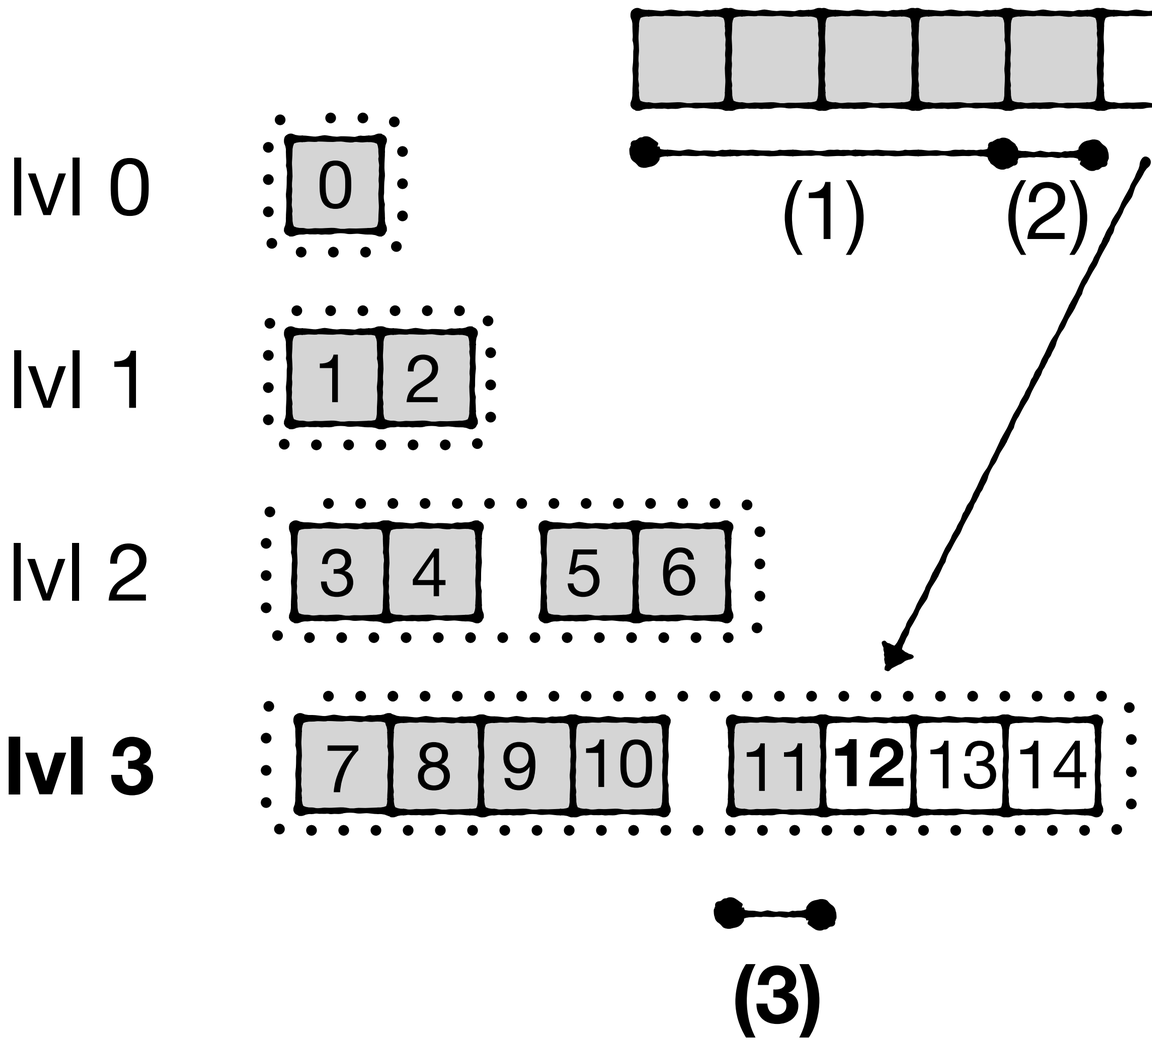

How to do steps 1 to 3 super fast?

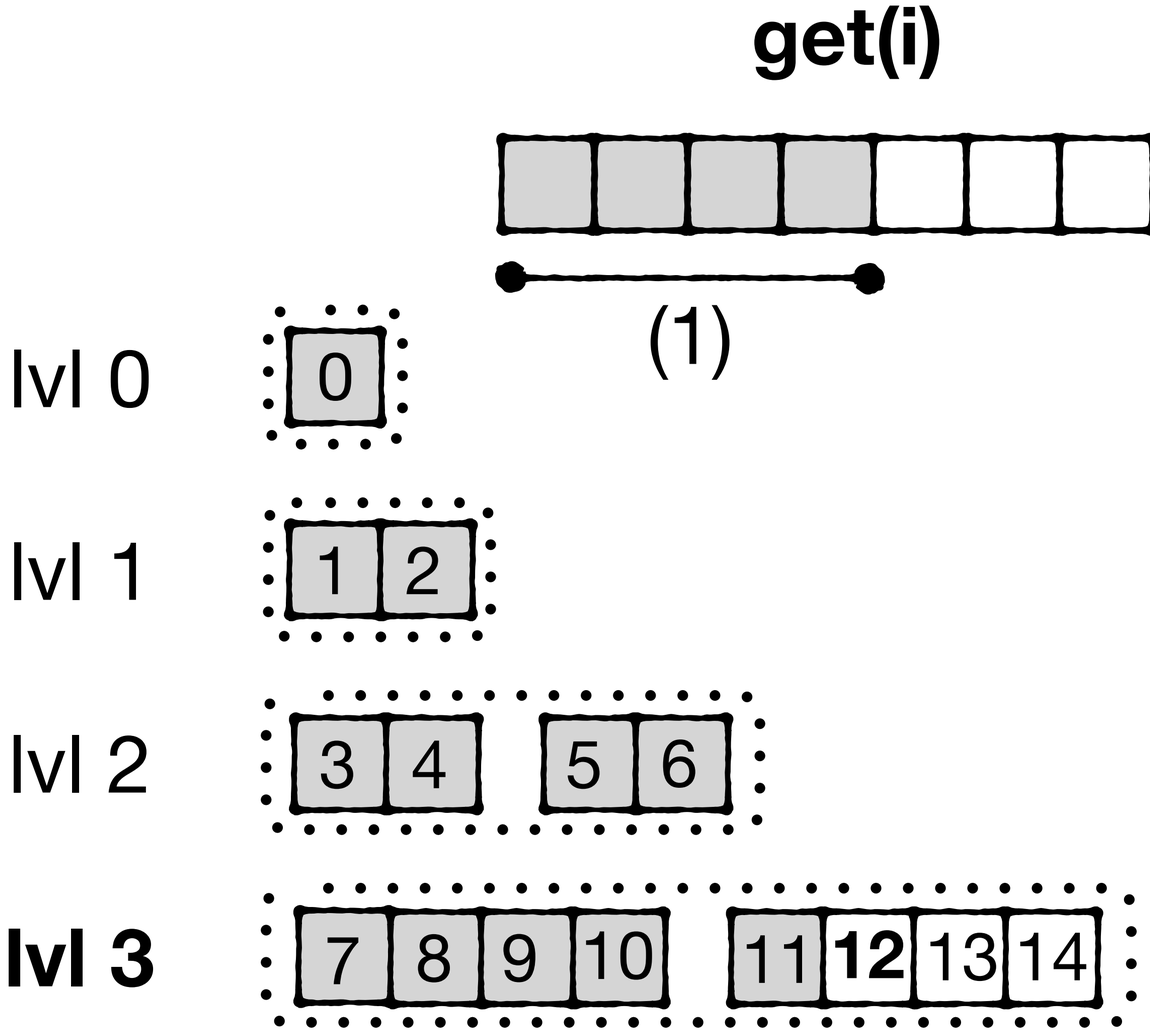
get(i)
(1)
lvl 0
0
lvl 1
1
2
lvl 2
3
4
5
6
lvl 3
7
8
9
10
11
12
13
14

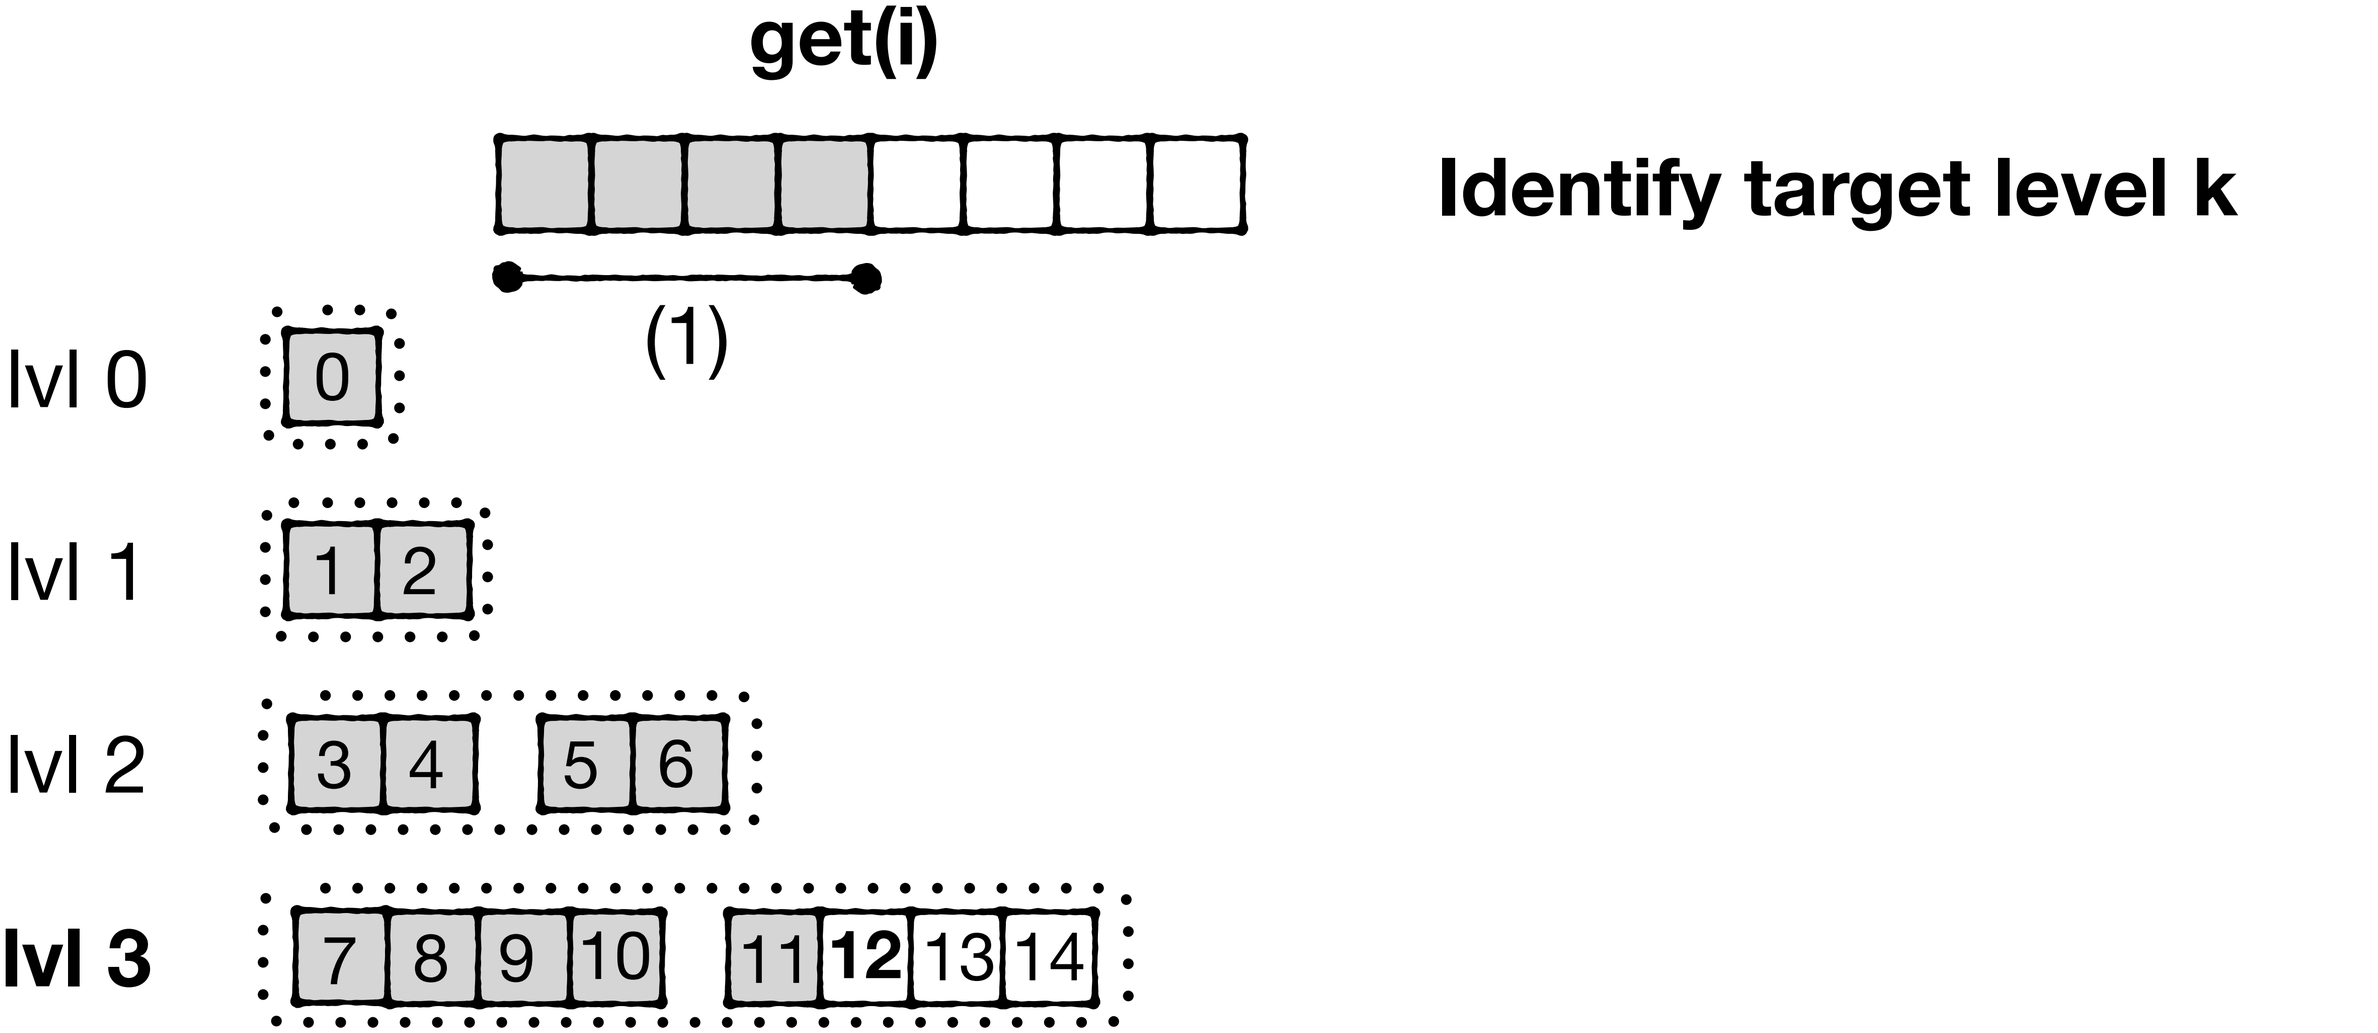

Identify target level k

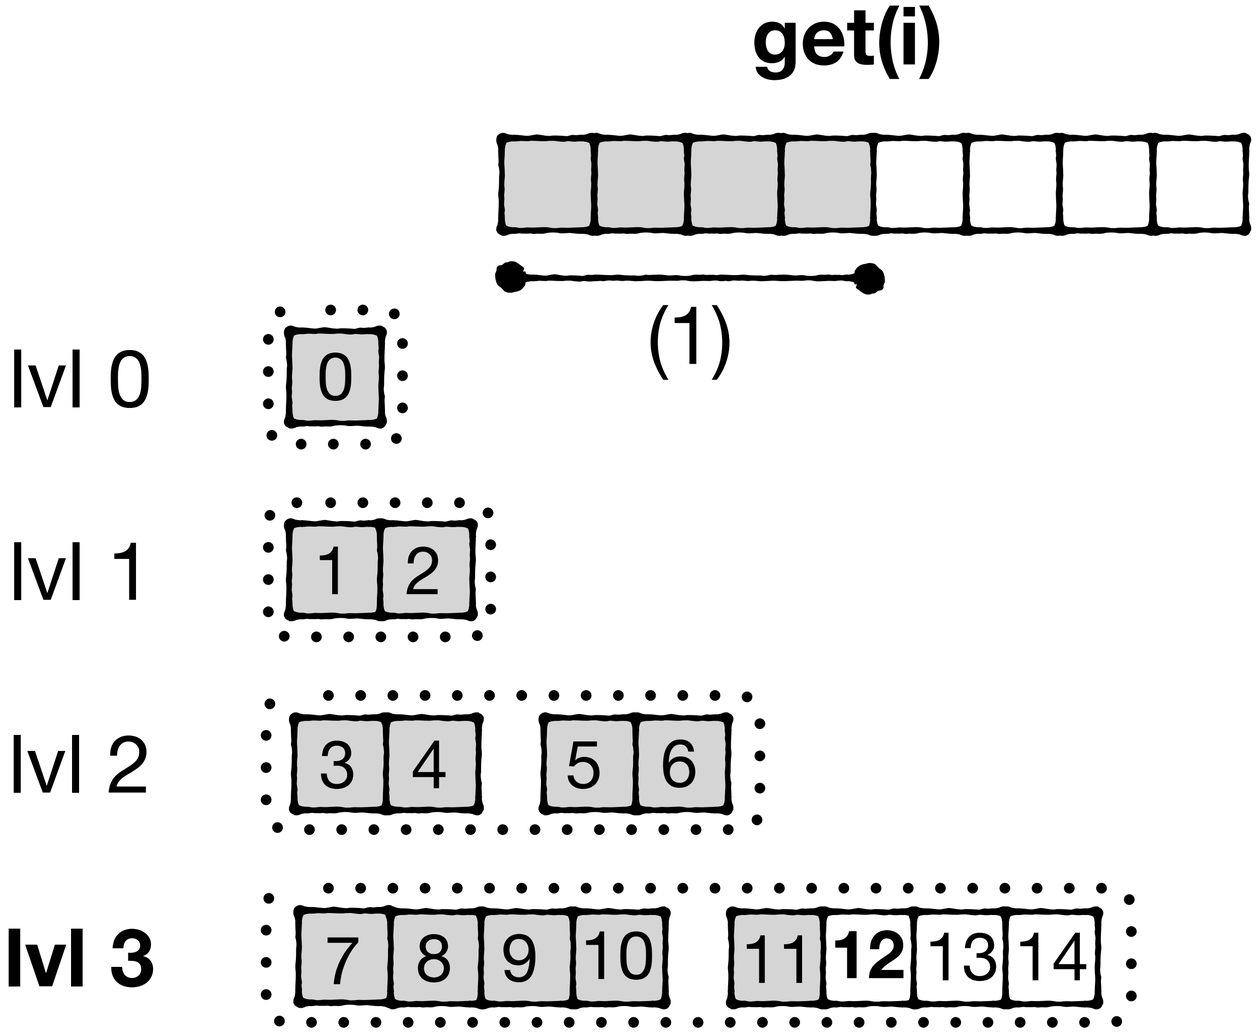

Identify target level k

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

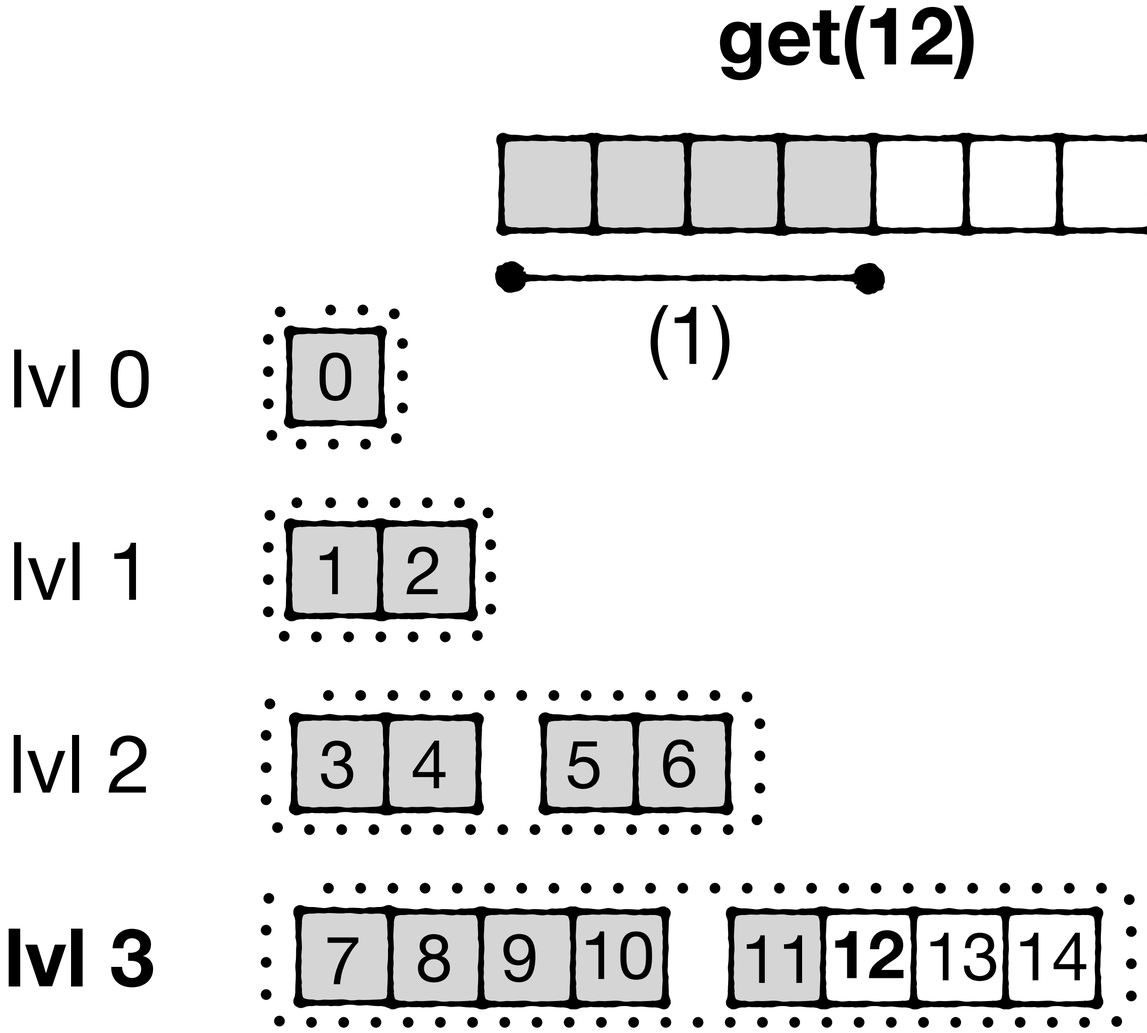

Identify target level k

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

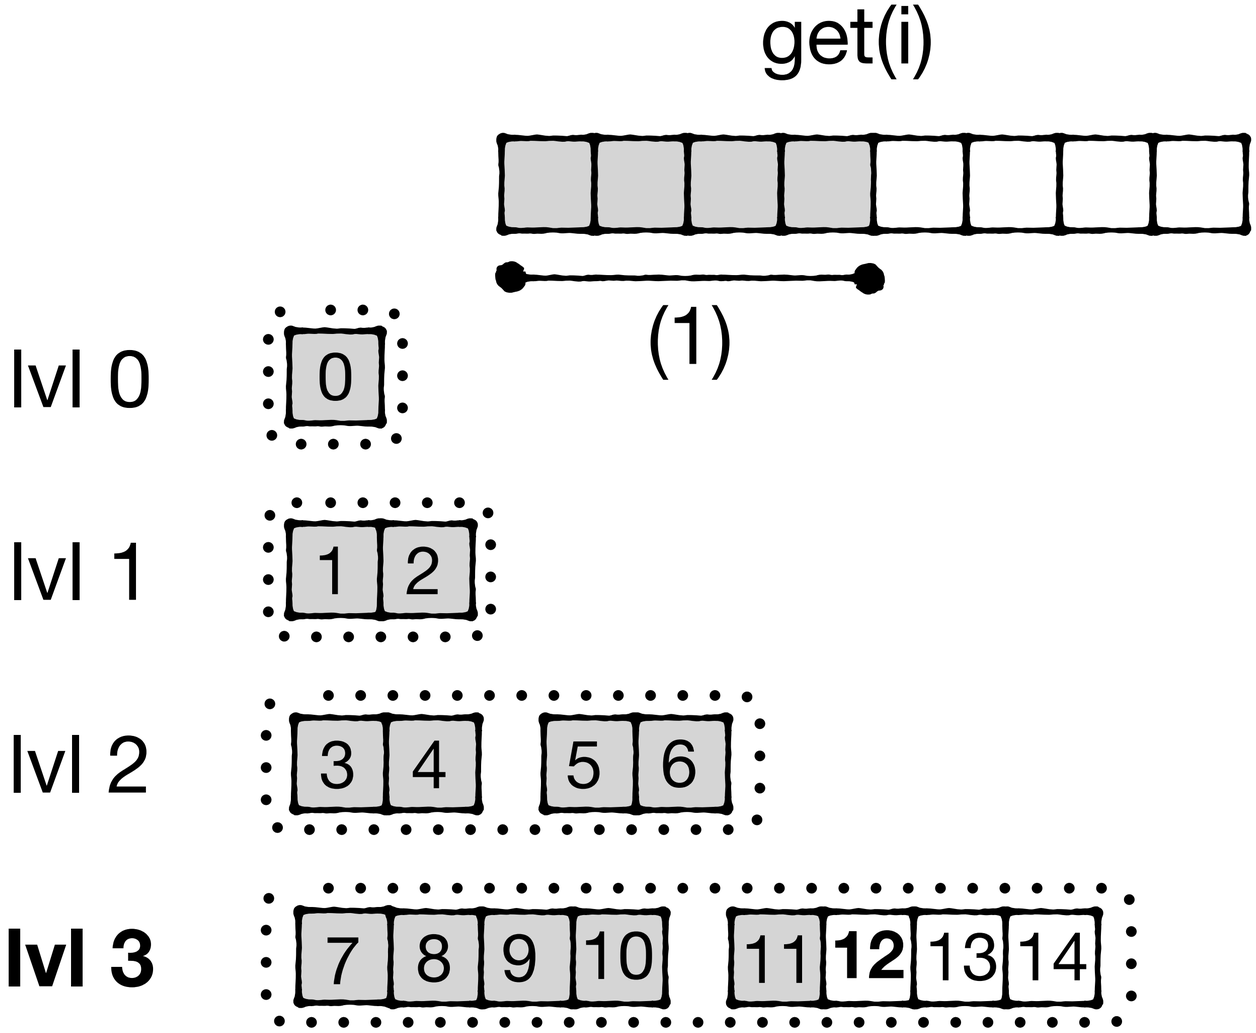

Identify target level k

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

↑

slow

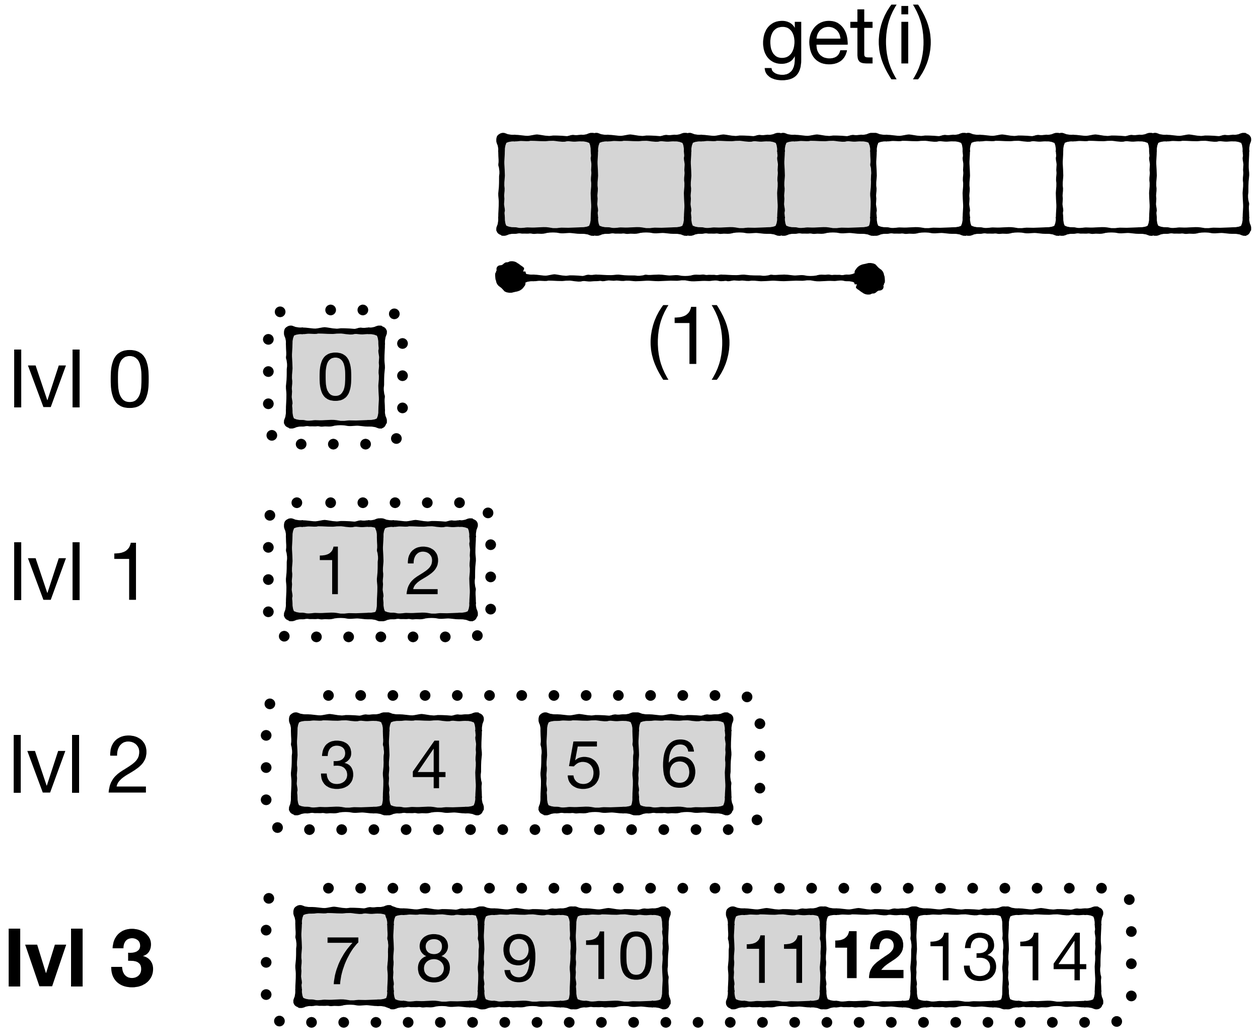

Identify target level k

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

Type casting - also slow

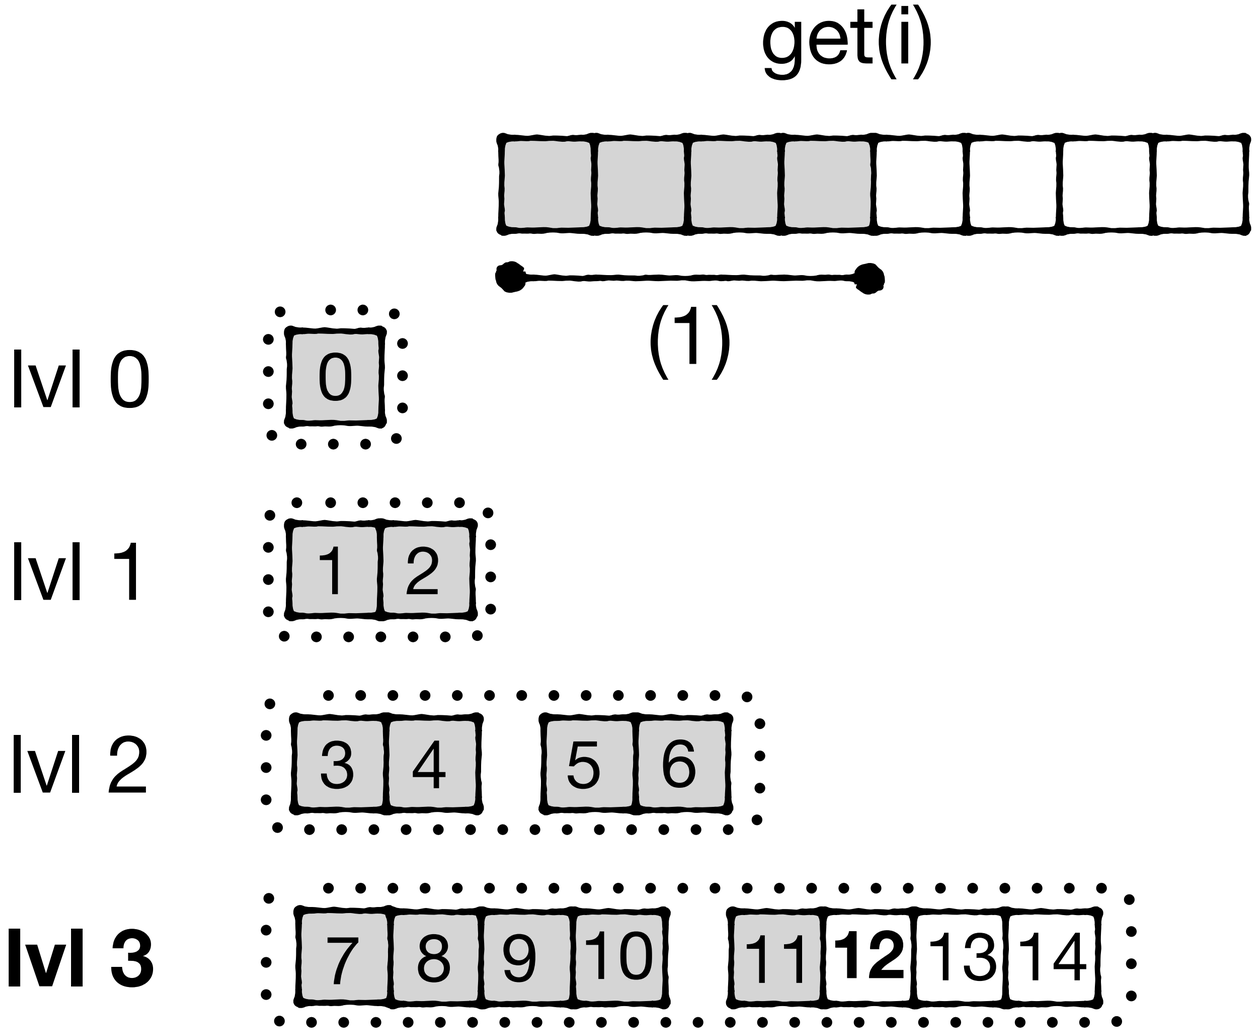

Identify target level k

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

Insight?

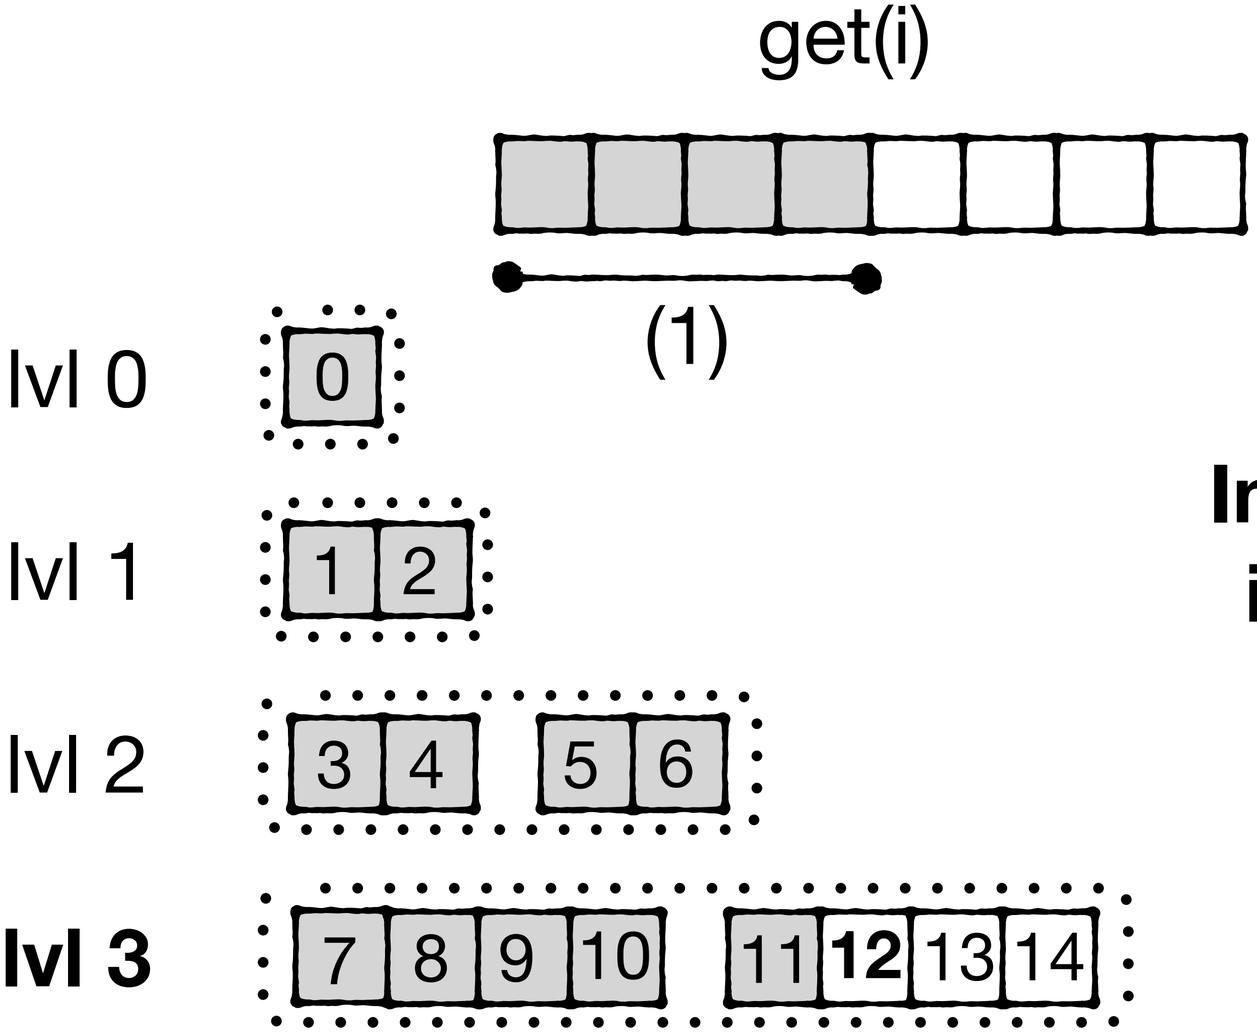

Identify target level k

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

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

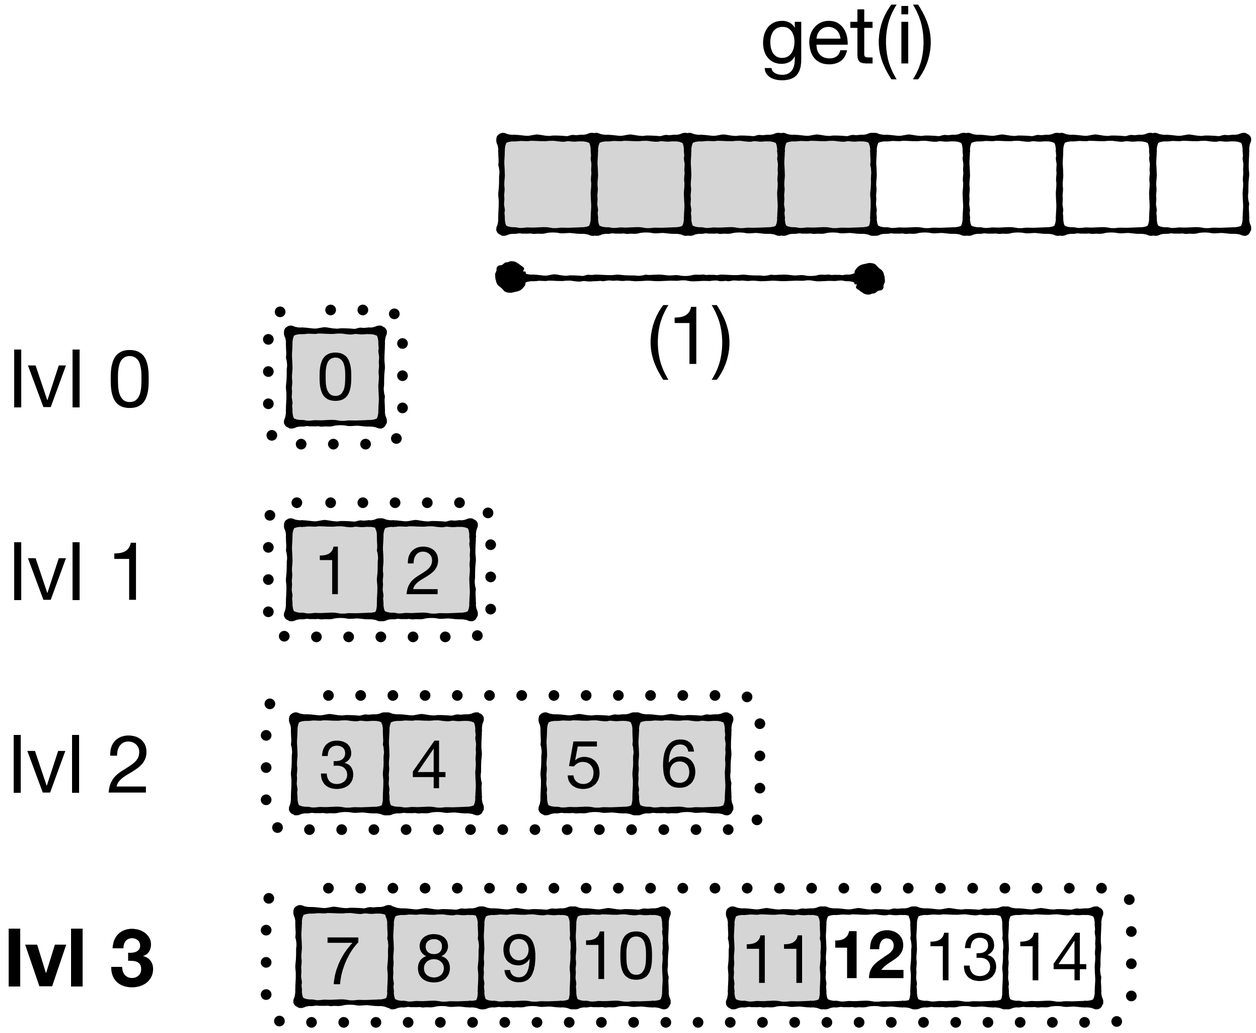

Identify target level k

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

= sizeof(i) - 1 - clz(i+1)

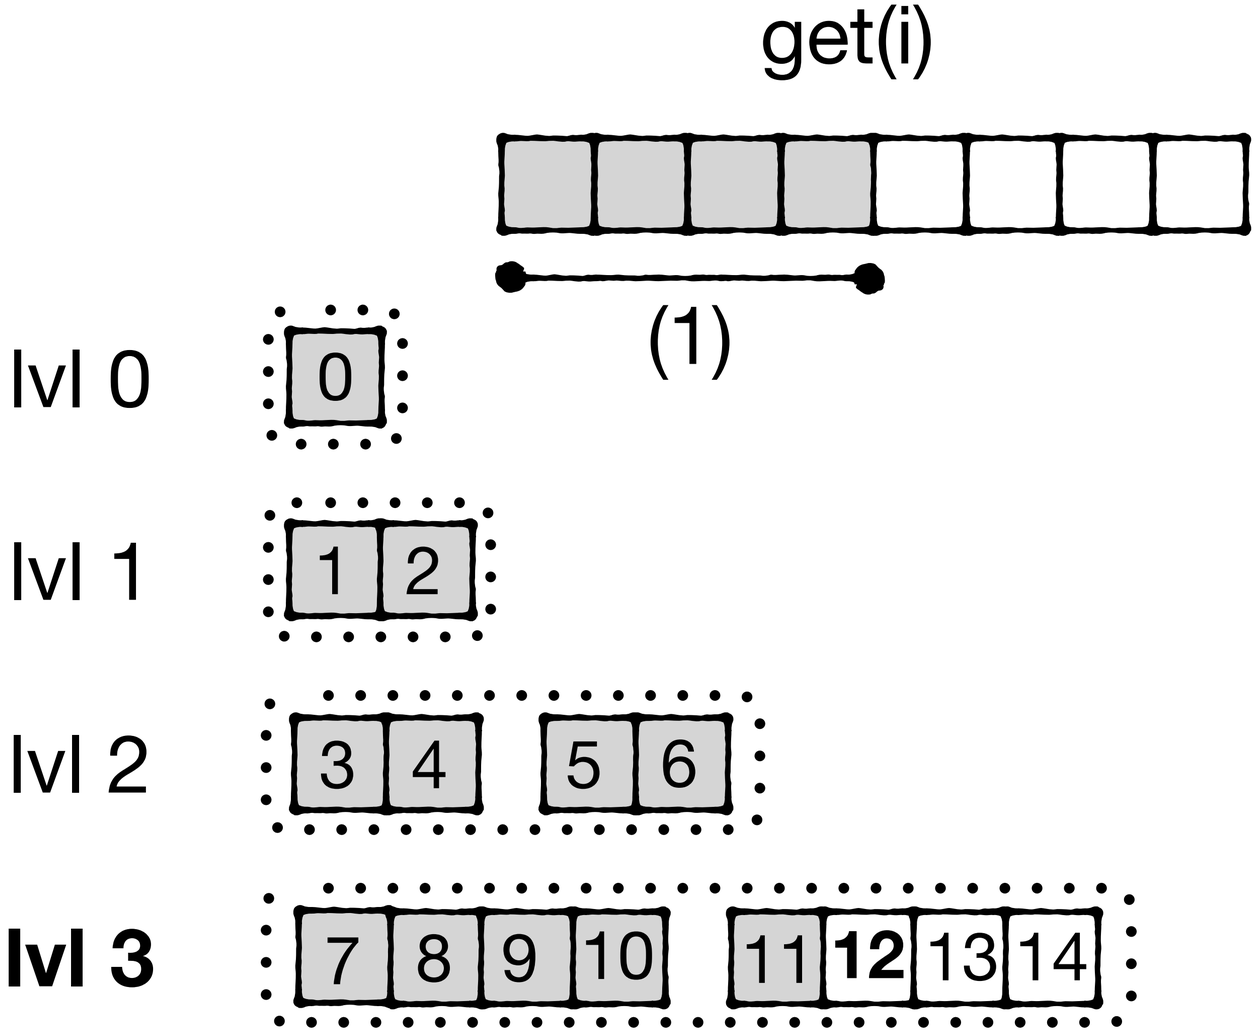

Identify target level k

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

= sizeof(i) - 1 - clz(i+1)

↑

Integer length in bits

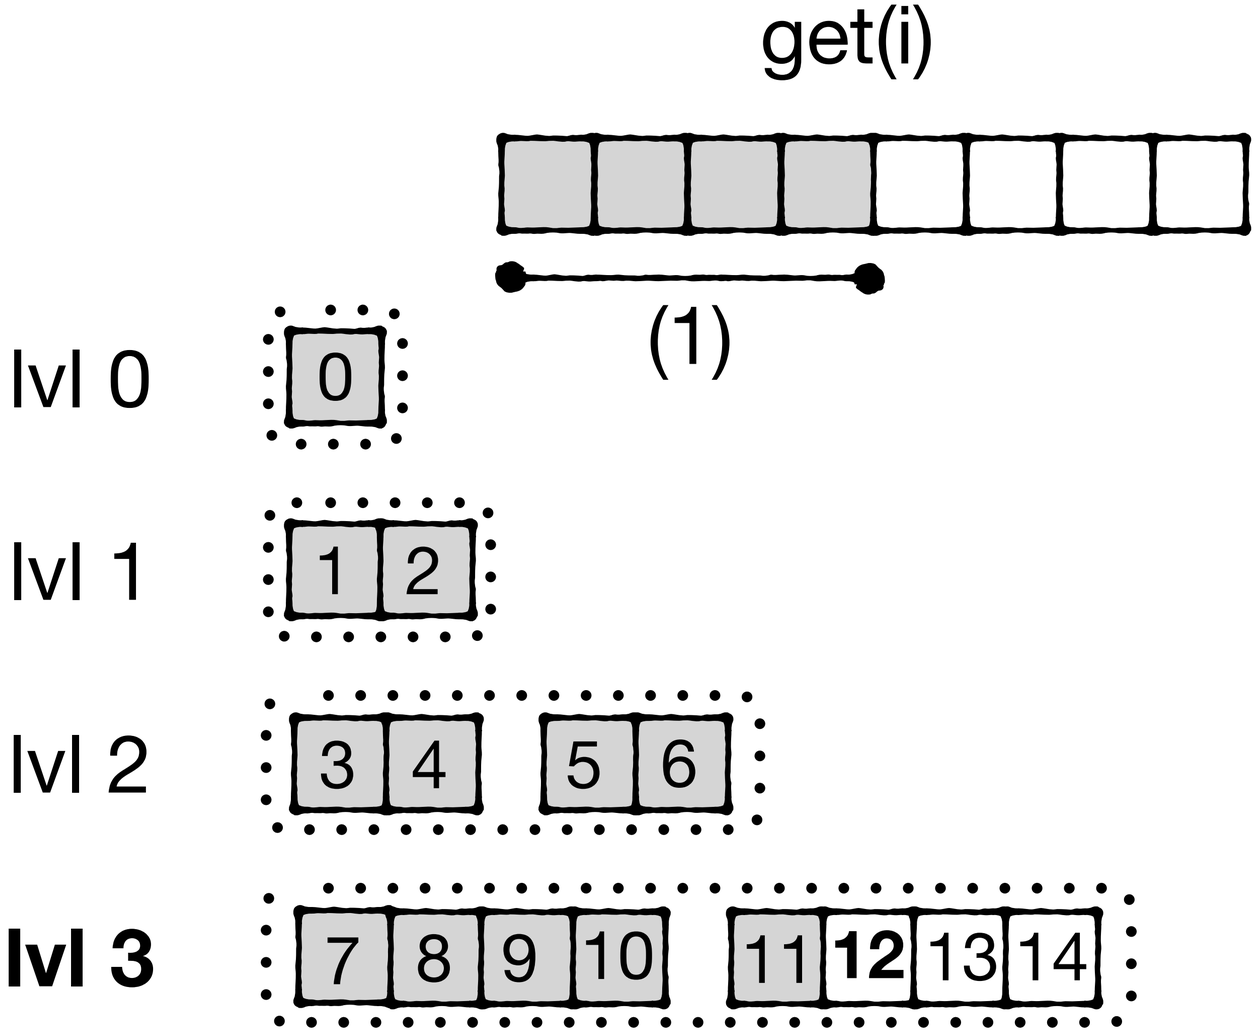

Identify target level k

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

= sizeof(i) - 1 - **clz(i+1)**

Specialized CPU command for # leading zeros

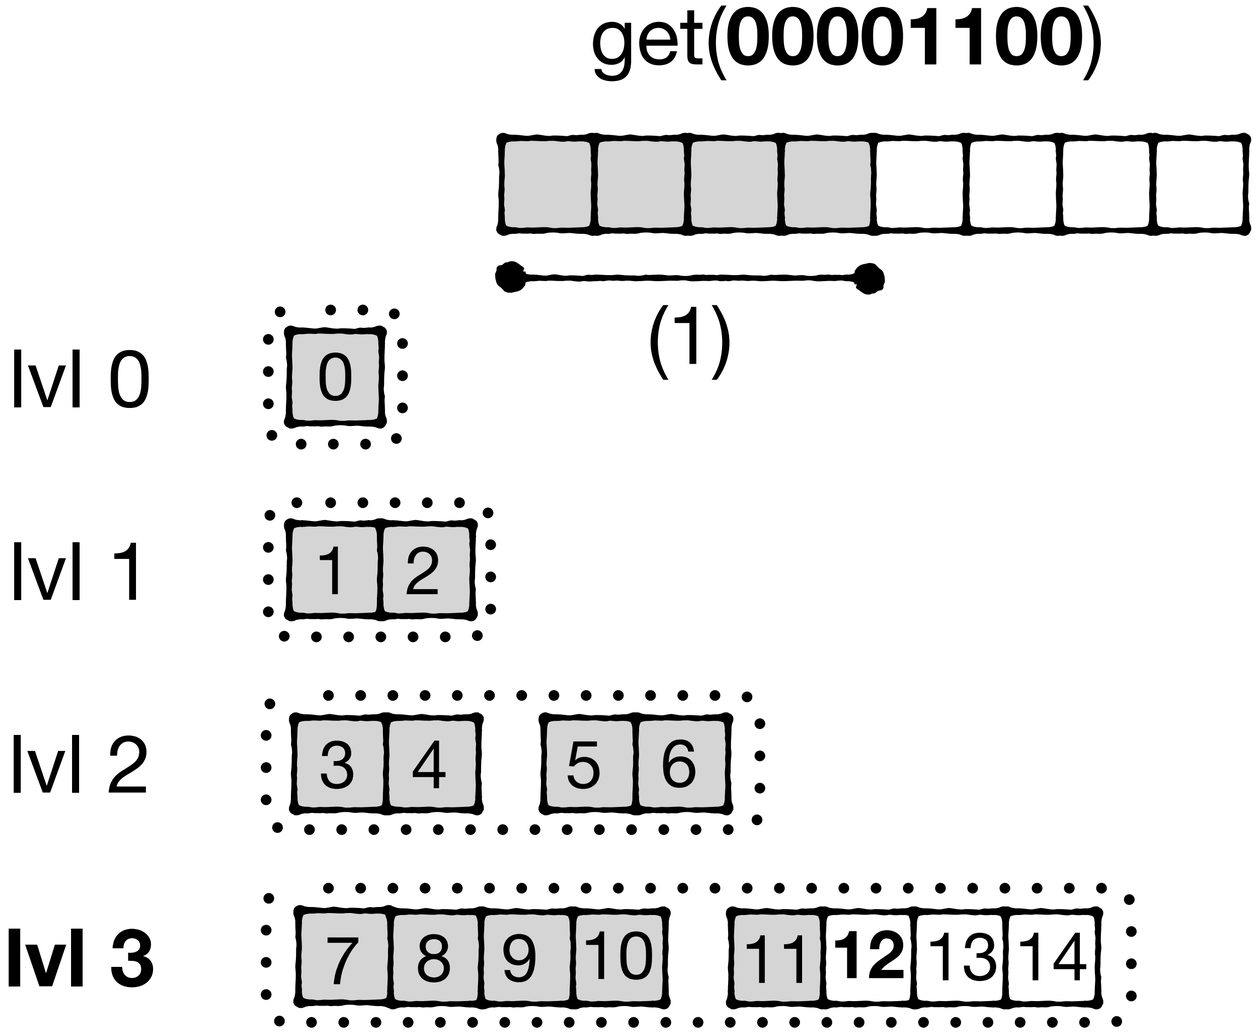

Identify target level k

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

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

8 - 1 - 4 = 3

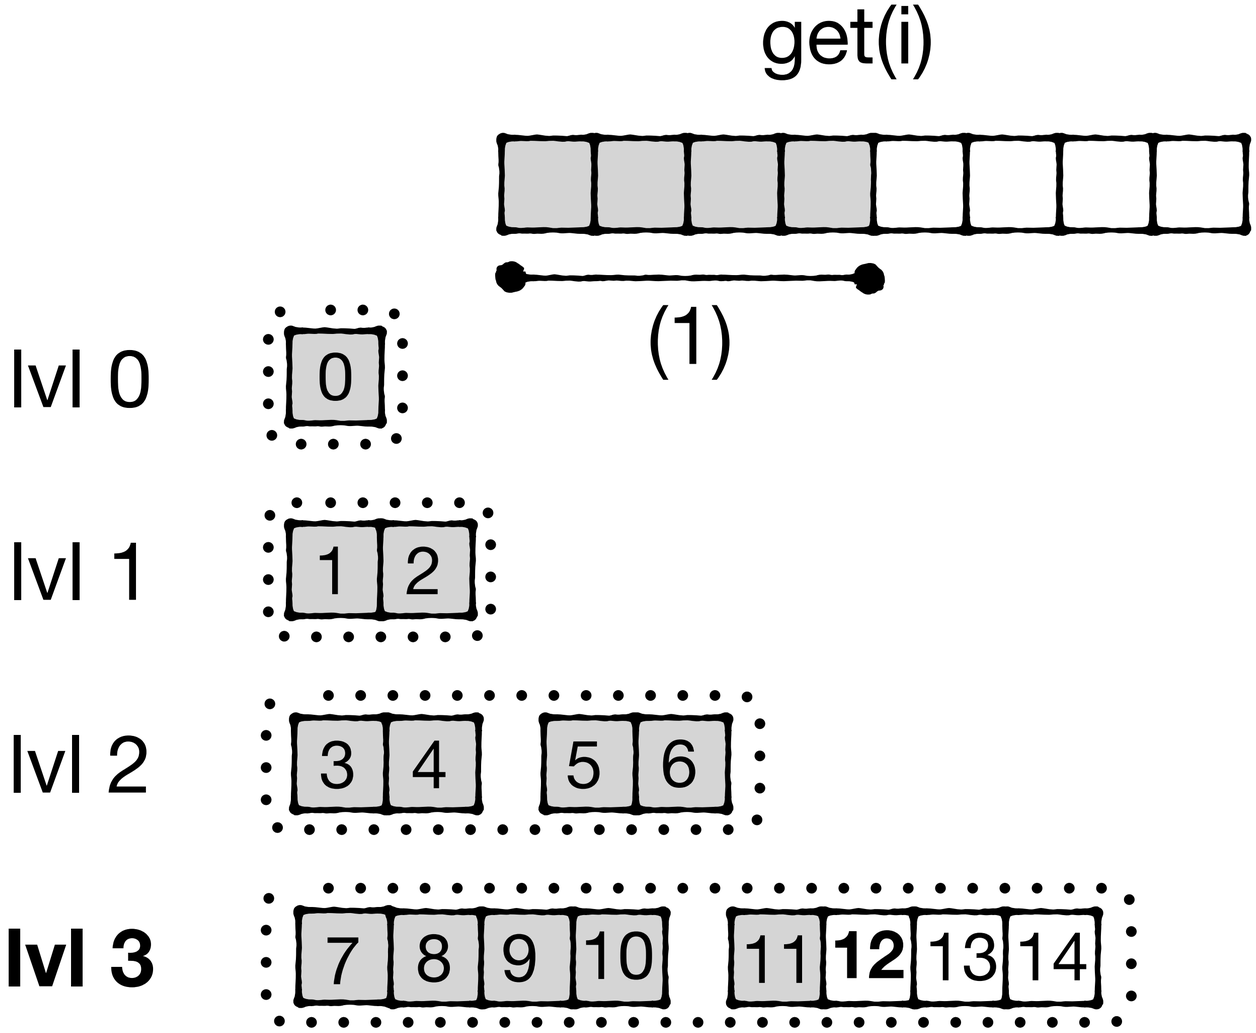

Identify target level k

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

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

≈1 ns rather than ≈7ns

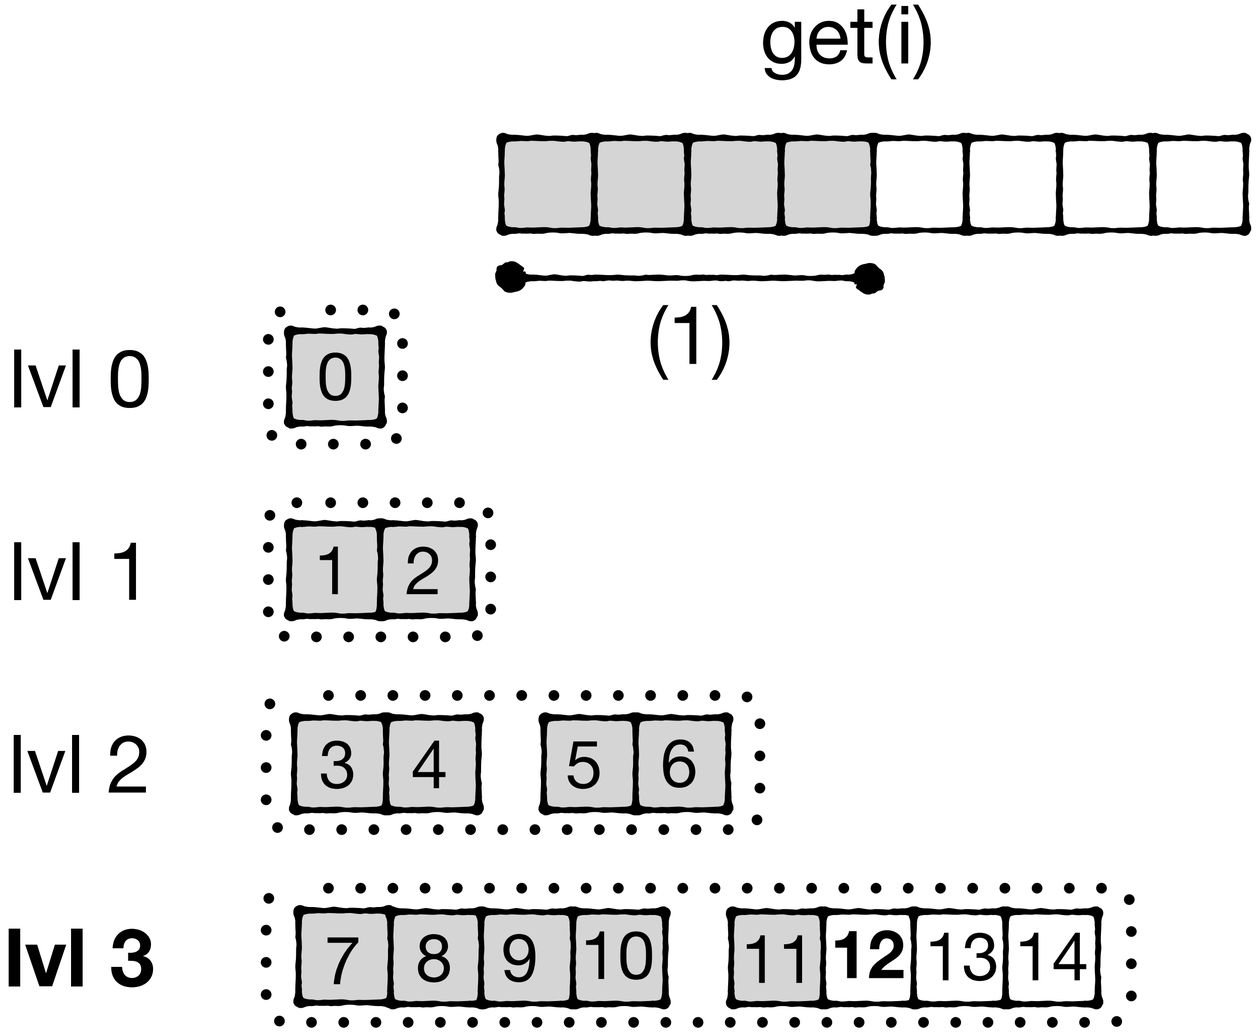

 # blocks in levels 0 to k-1?

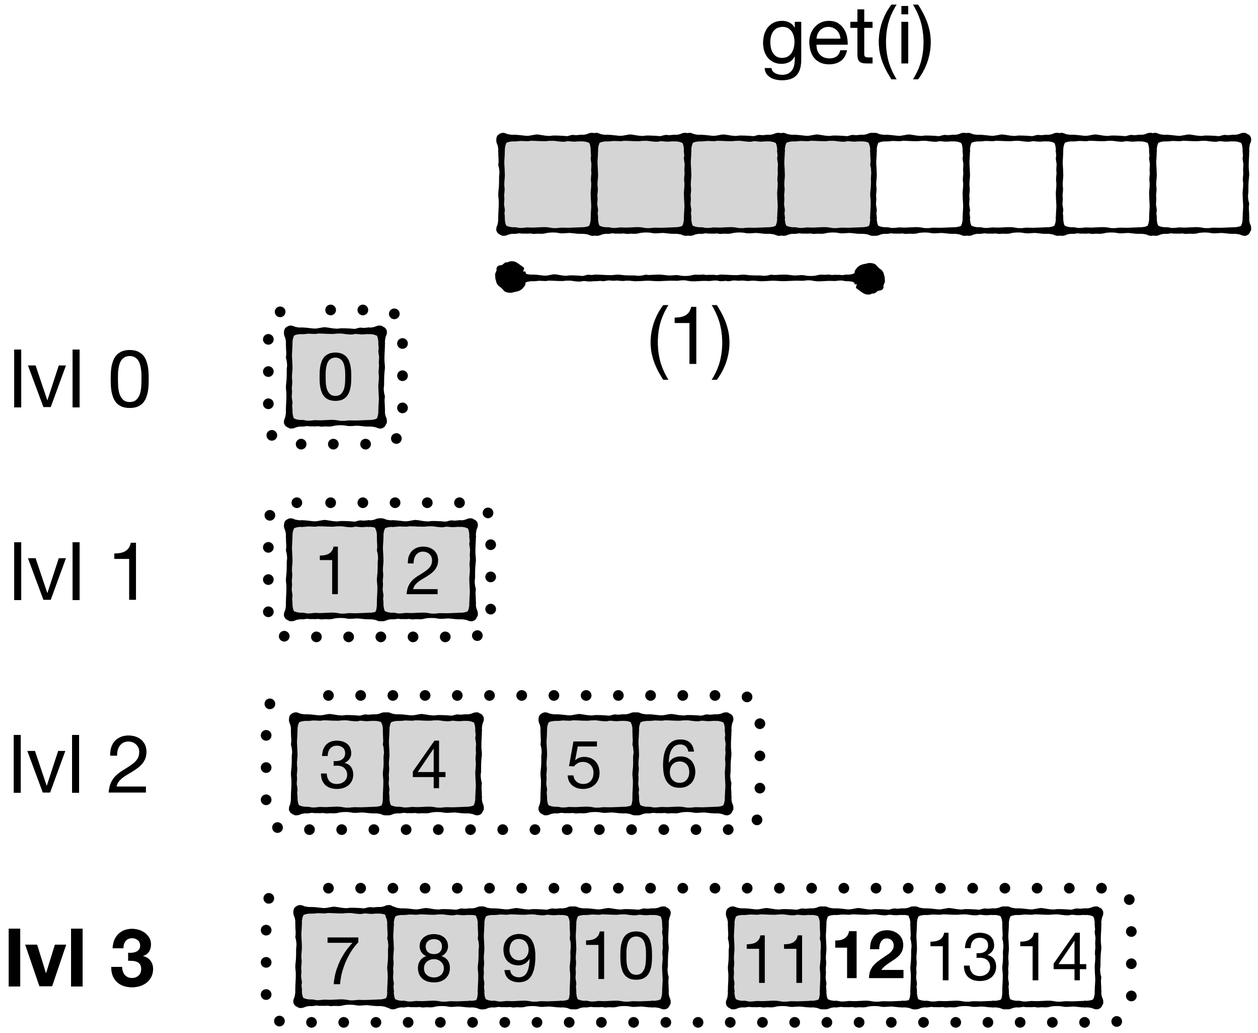

 # blocks in levels 0 to k-1

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

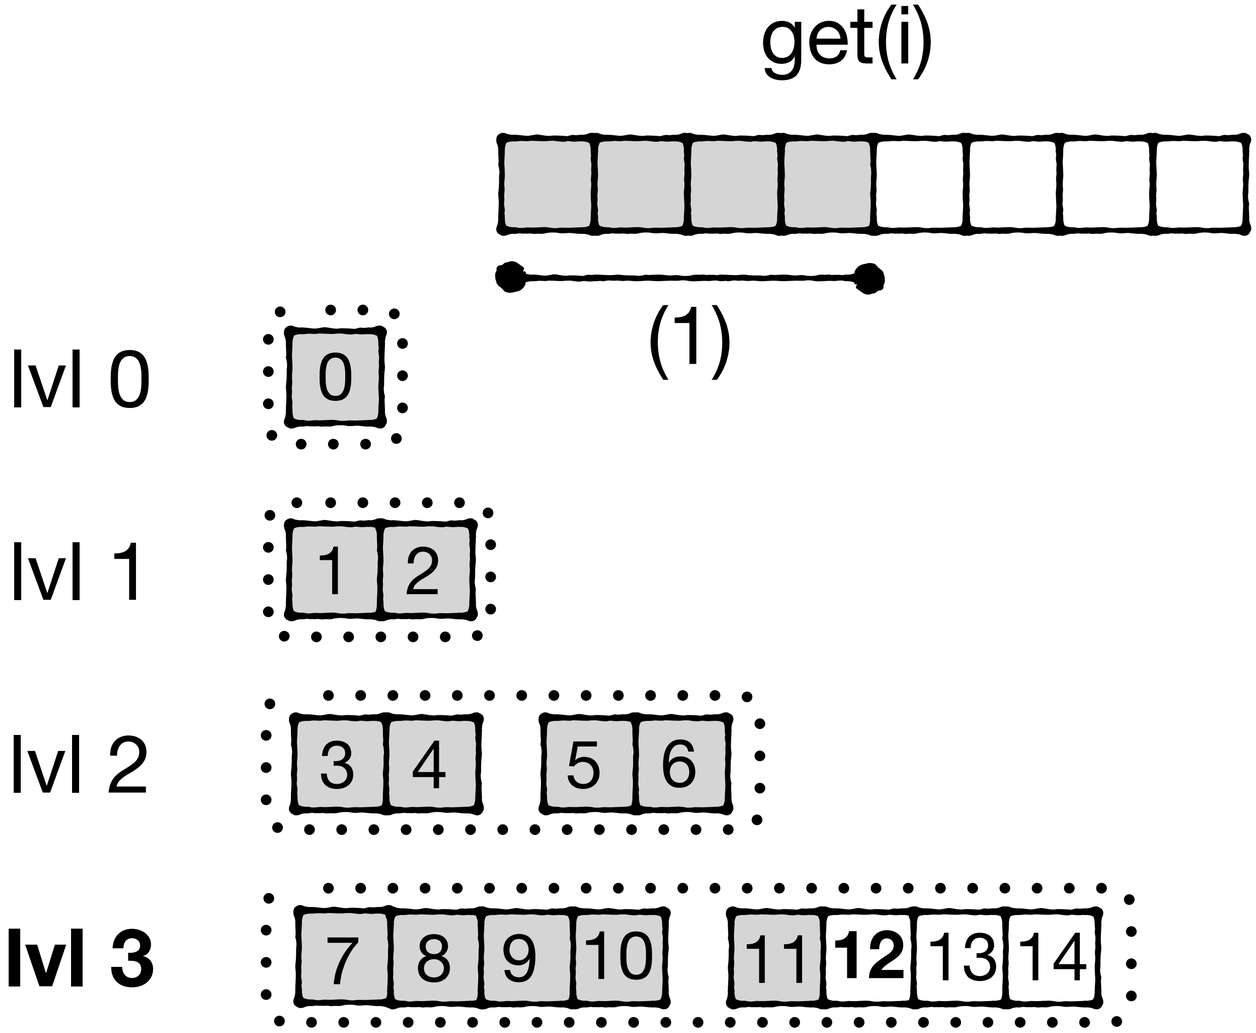

 # blocks in levels 0 to k-1

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

Original paper gets this wrong, fixed credit to Hyuhng Min

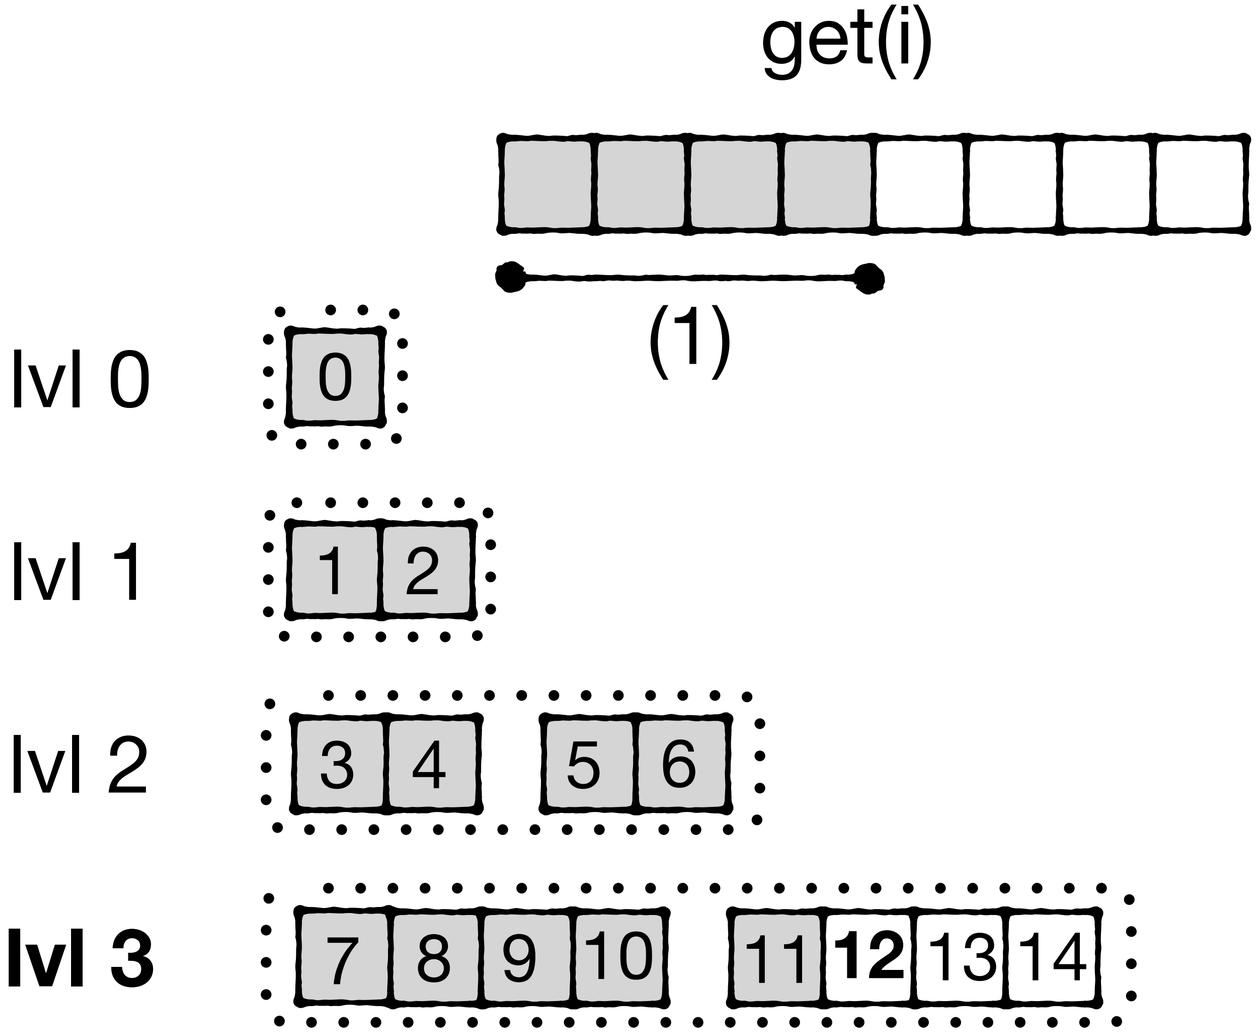

blocks in levels 0 to k-1

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

Intuition: number of new data blocks grows every other level

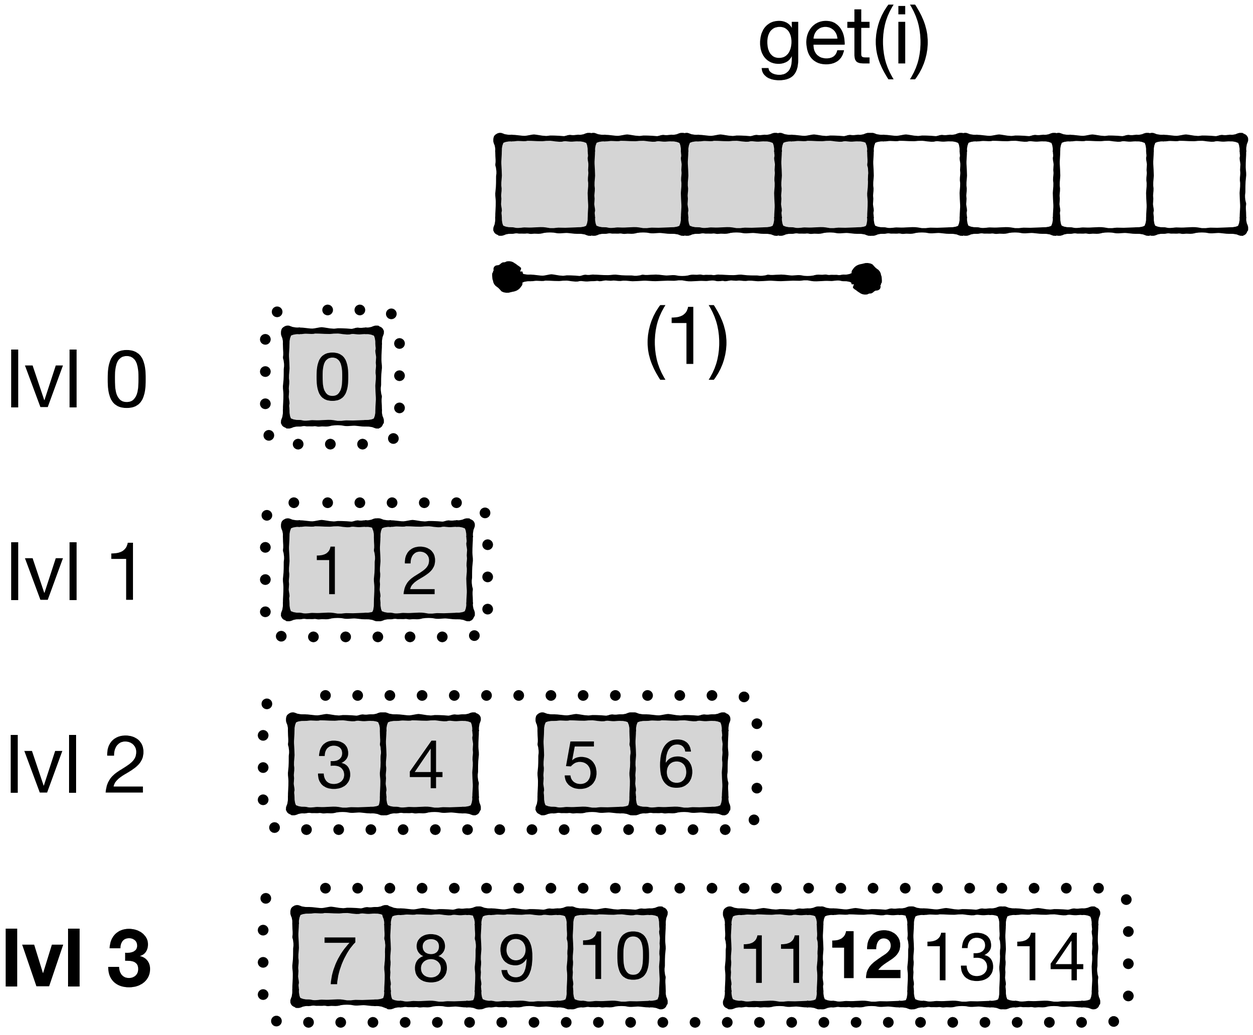

blocks in levels 0 to k-1

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

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

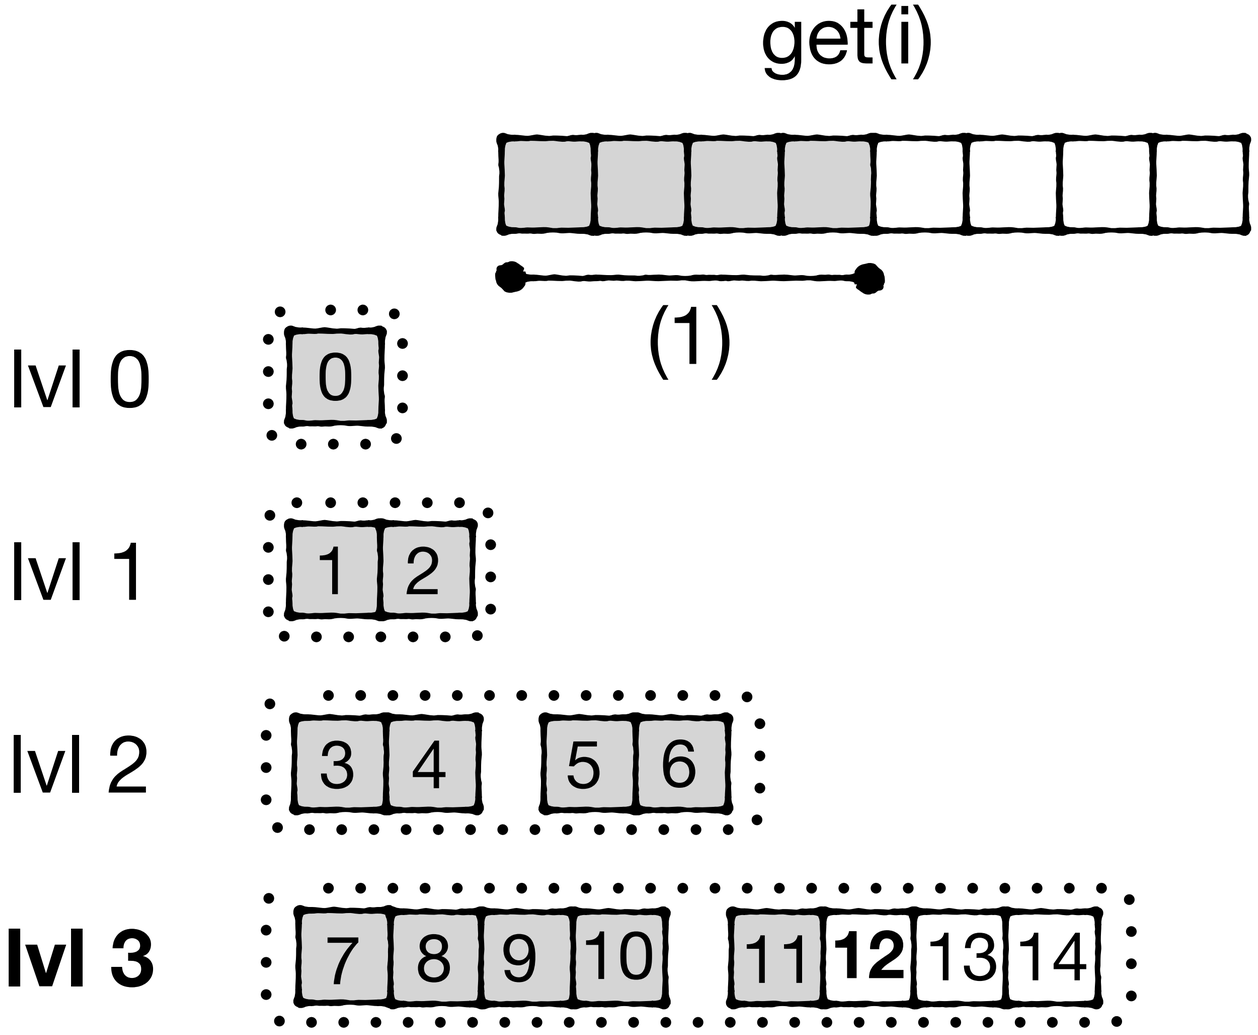

blocks in levels 0 to k-1

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

Integer division is slow

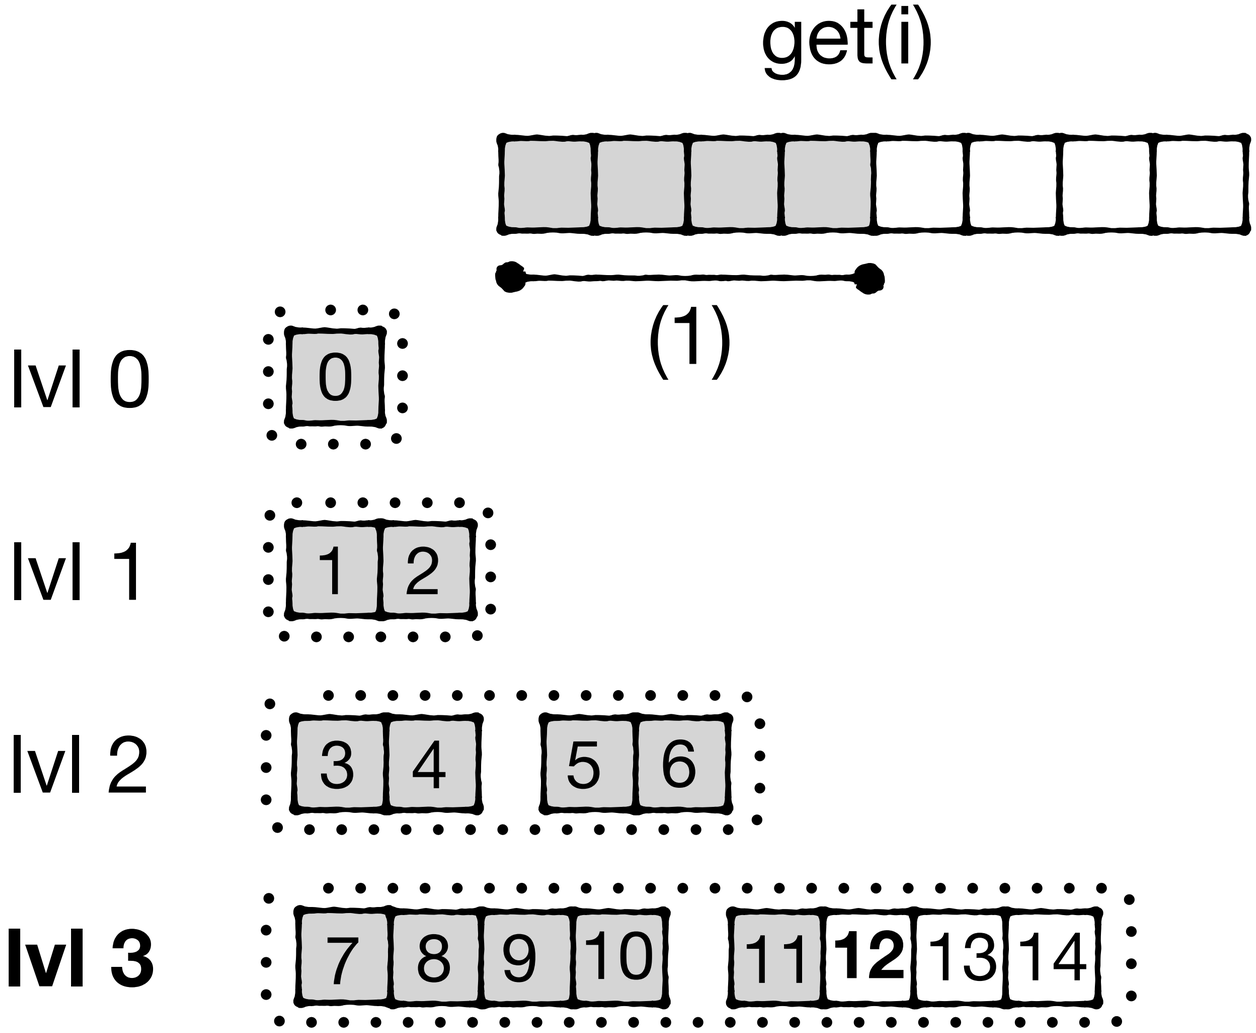

blocks in levels 0 to k-1

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

Power is slow

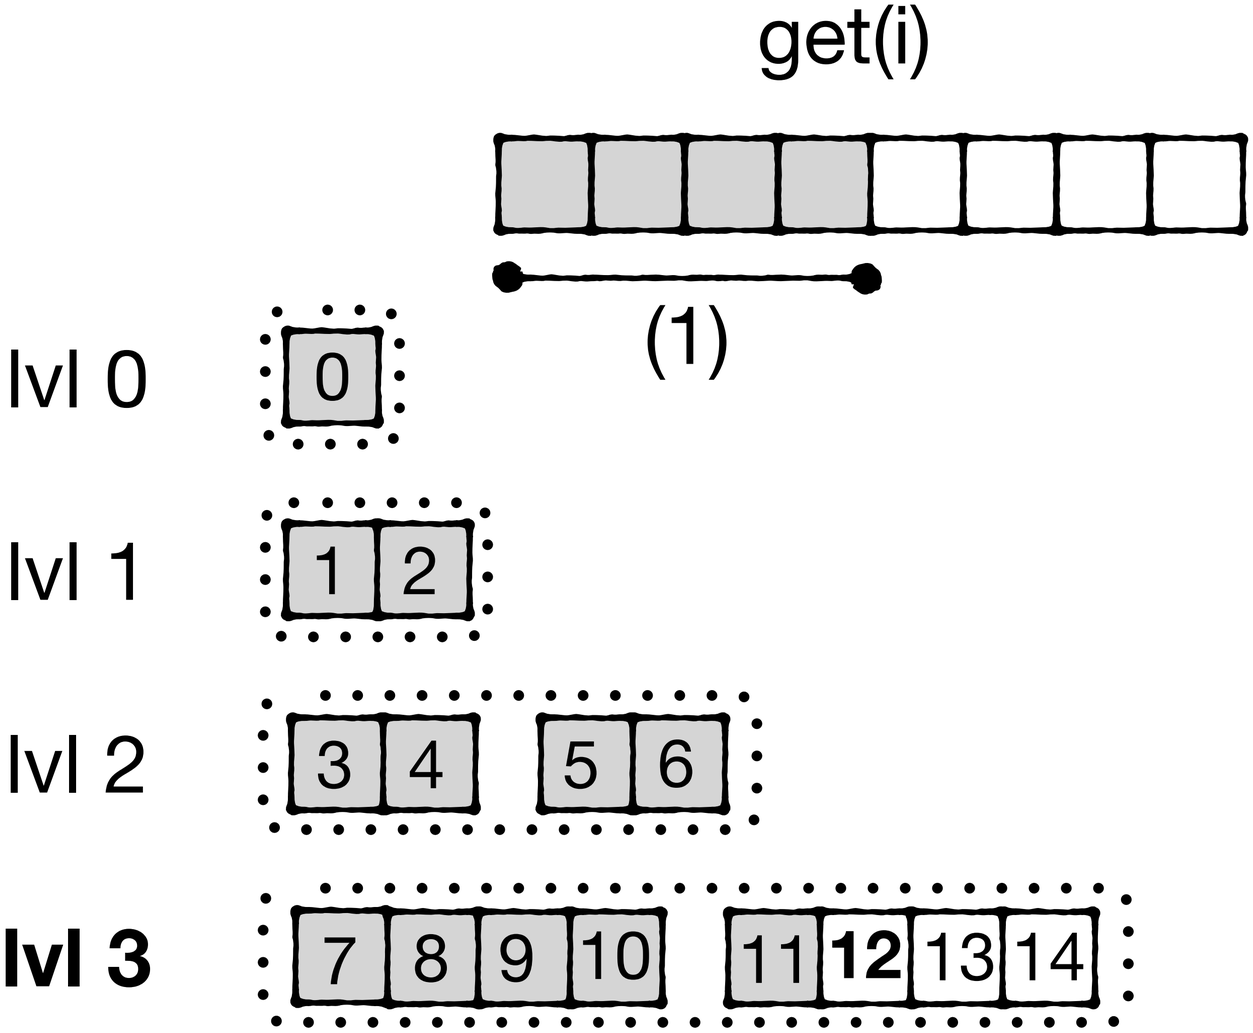

blocks in levels 0 to k-1

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

How to speed up?

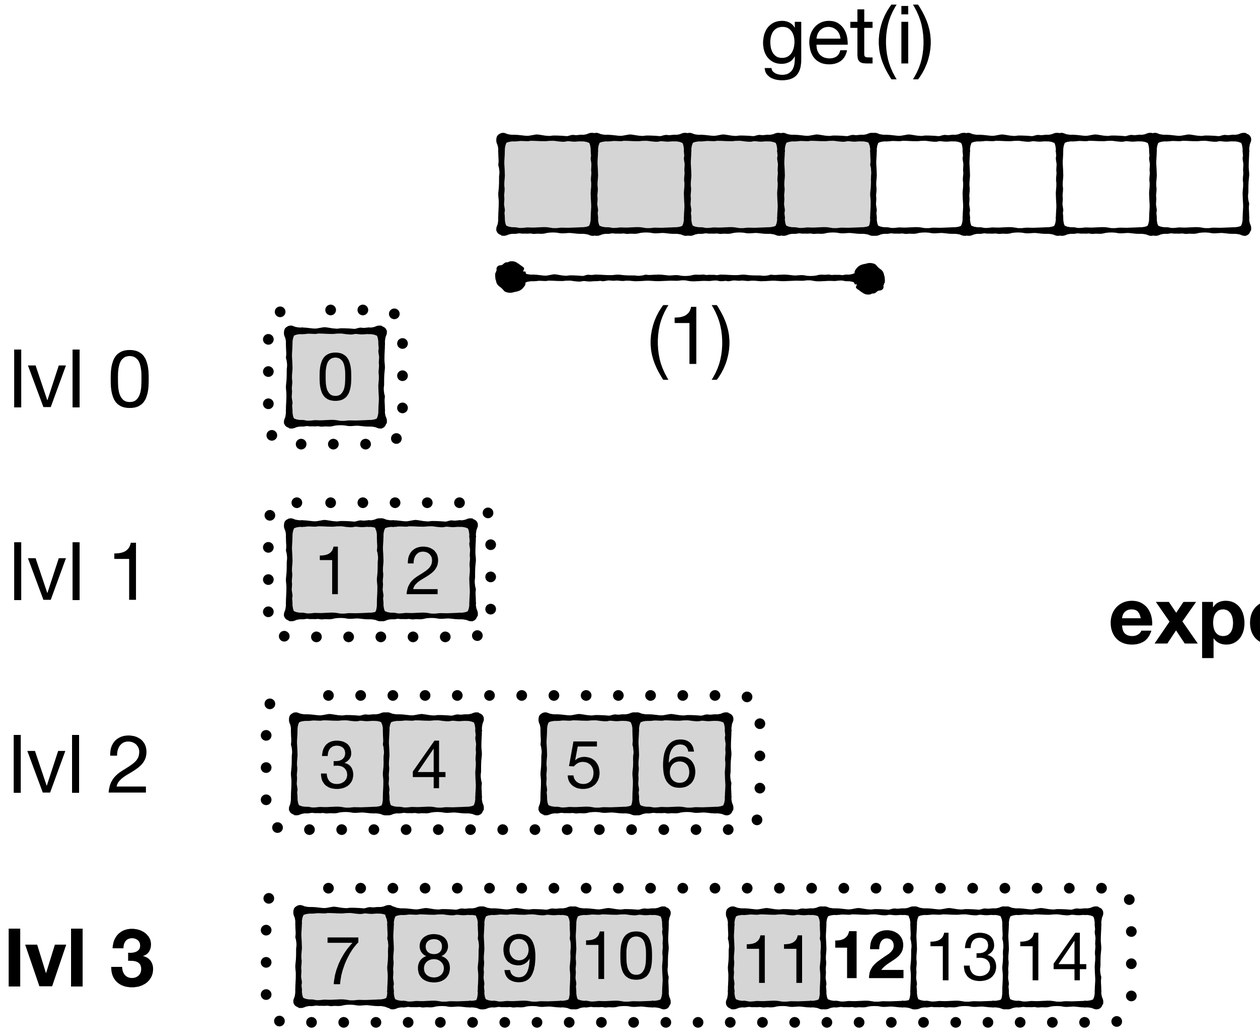

blocks in levels 0 to k-1

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

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

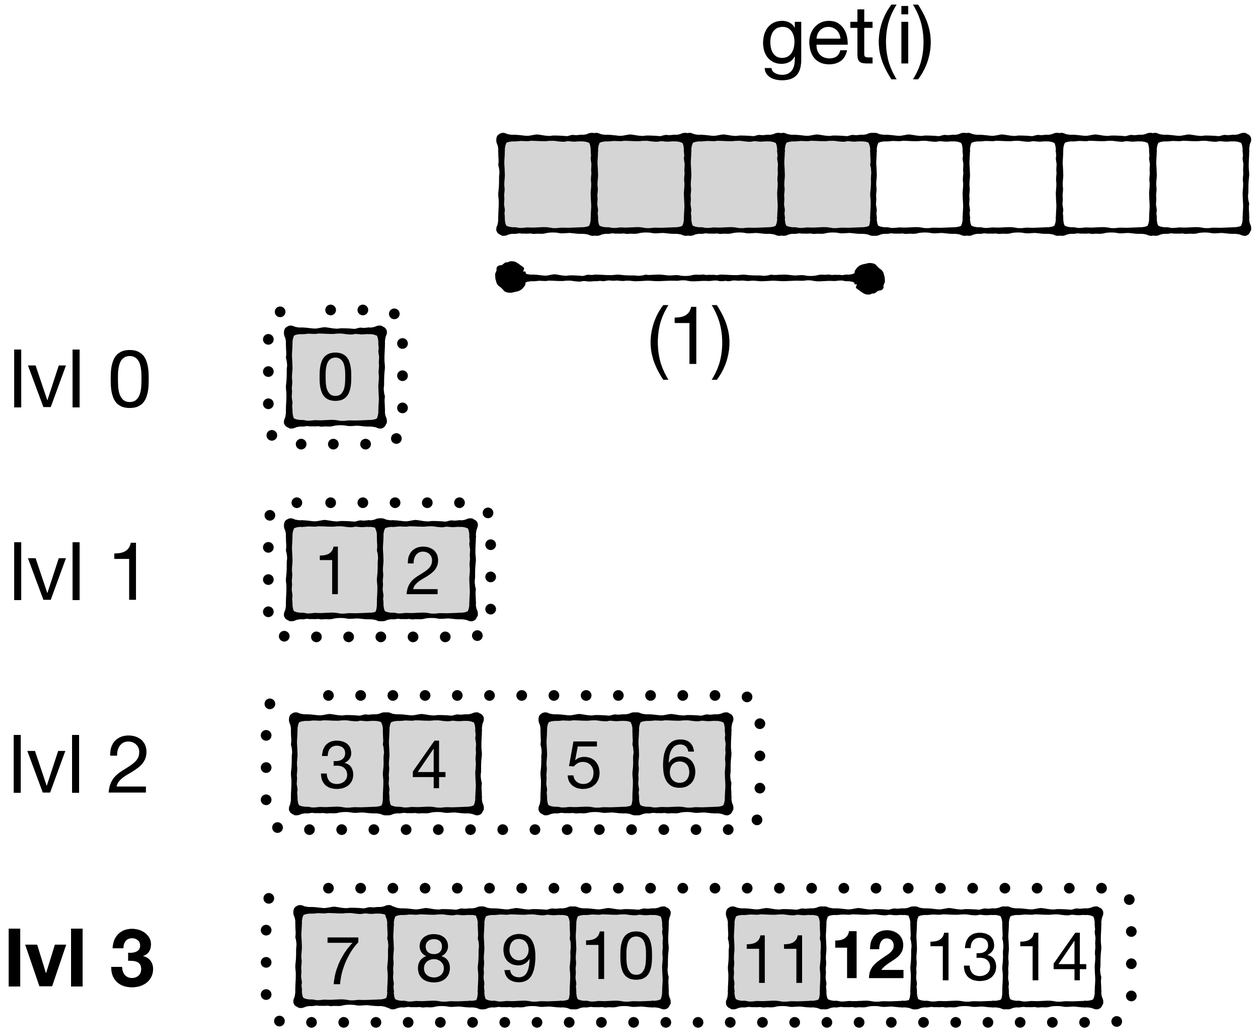

blocks in levels 0 to k-1

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

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

“and” with 1

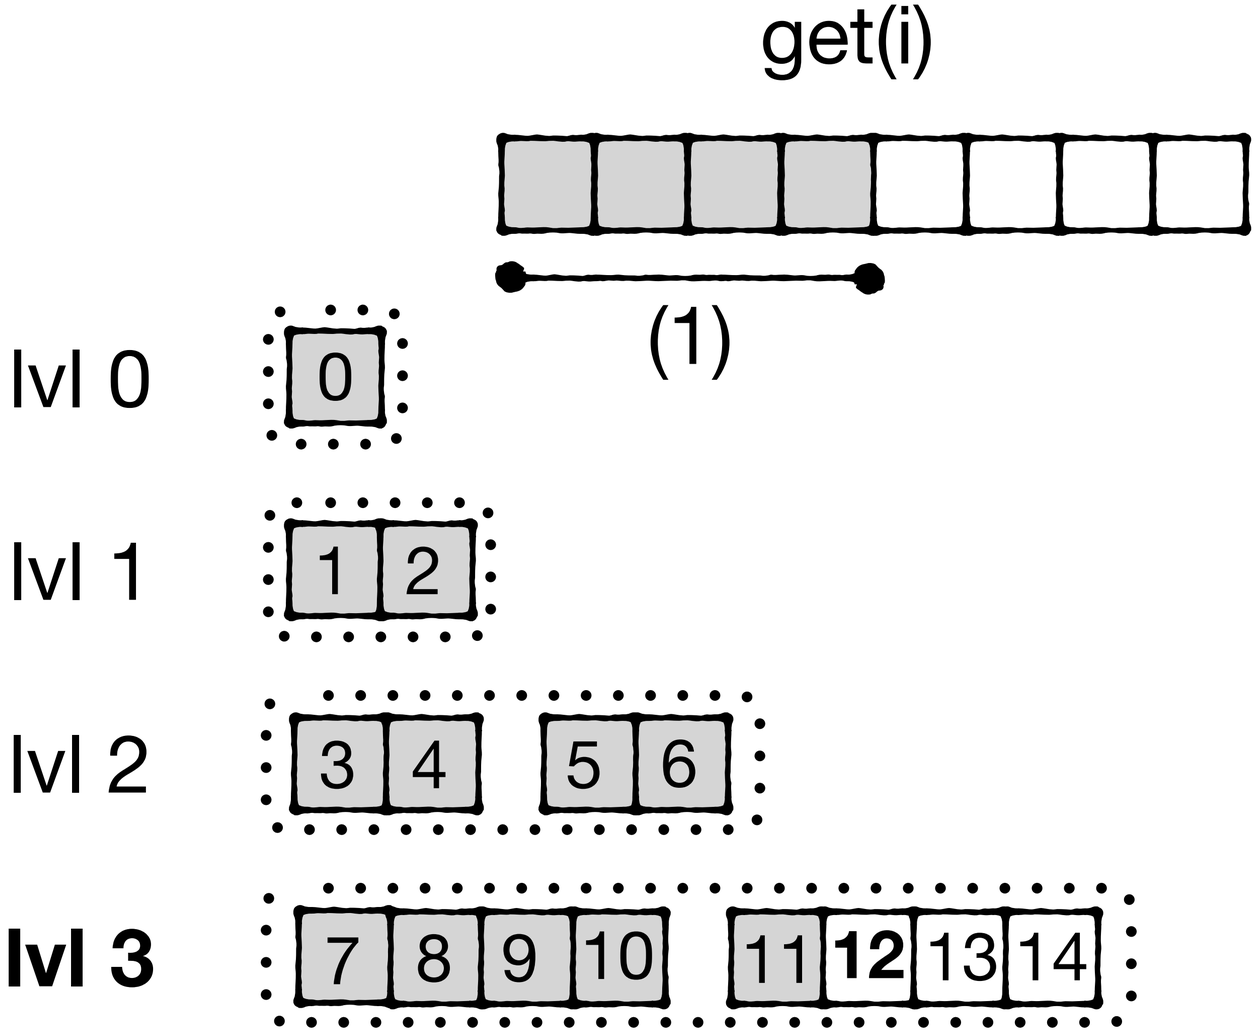

blocks in levels 0 to k-1

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

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

Shift by 1 bit to right

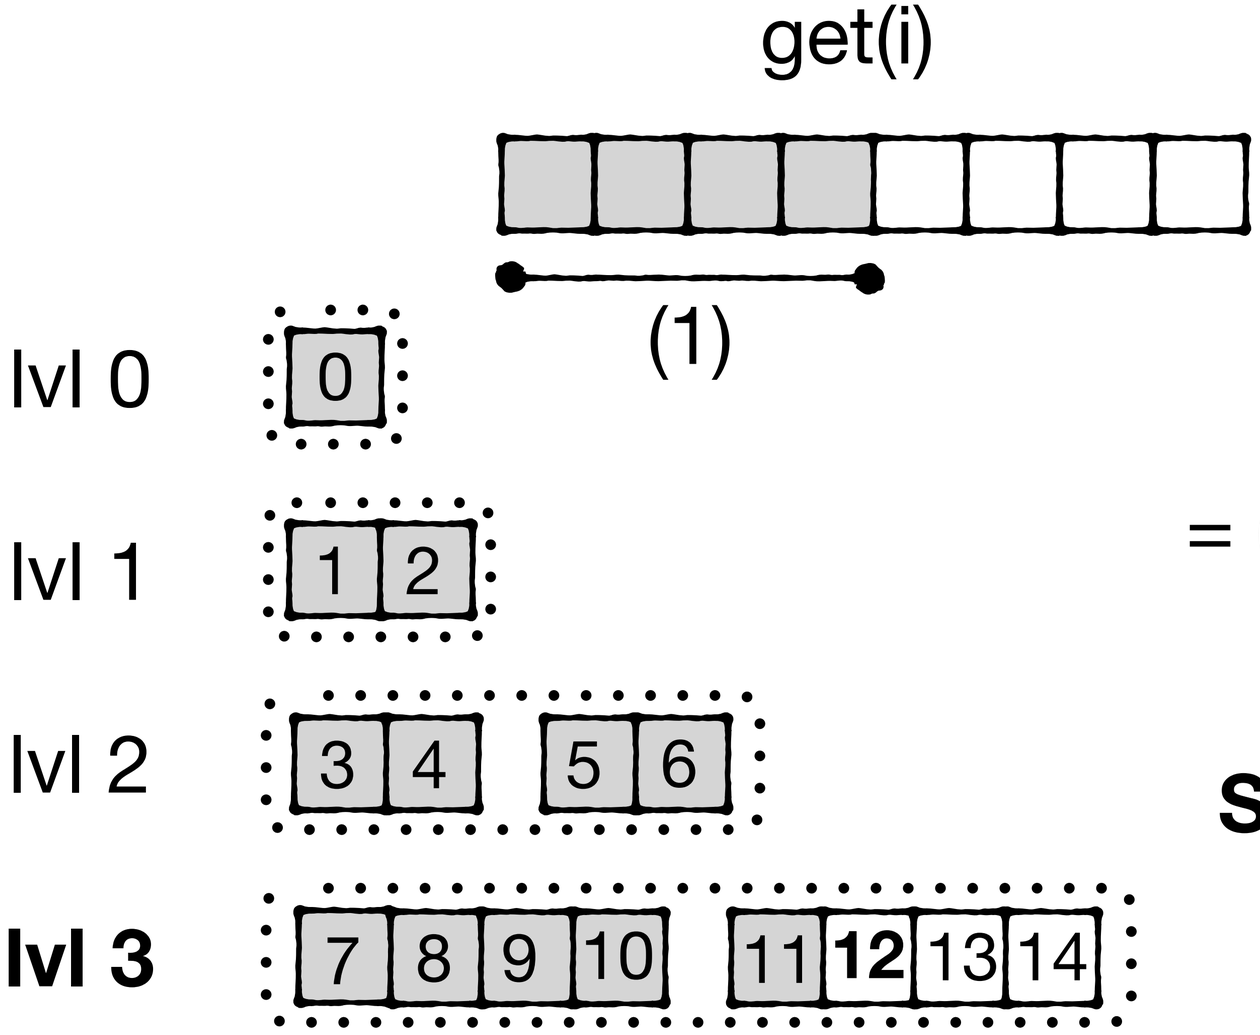

blocks in levels 0 to k-1

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

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

Shift to left

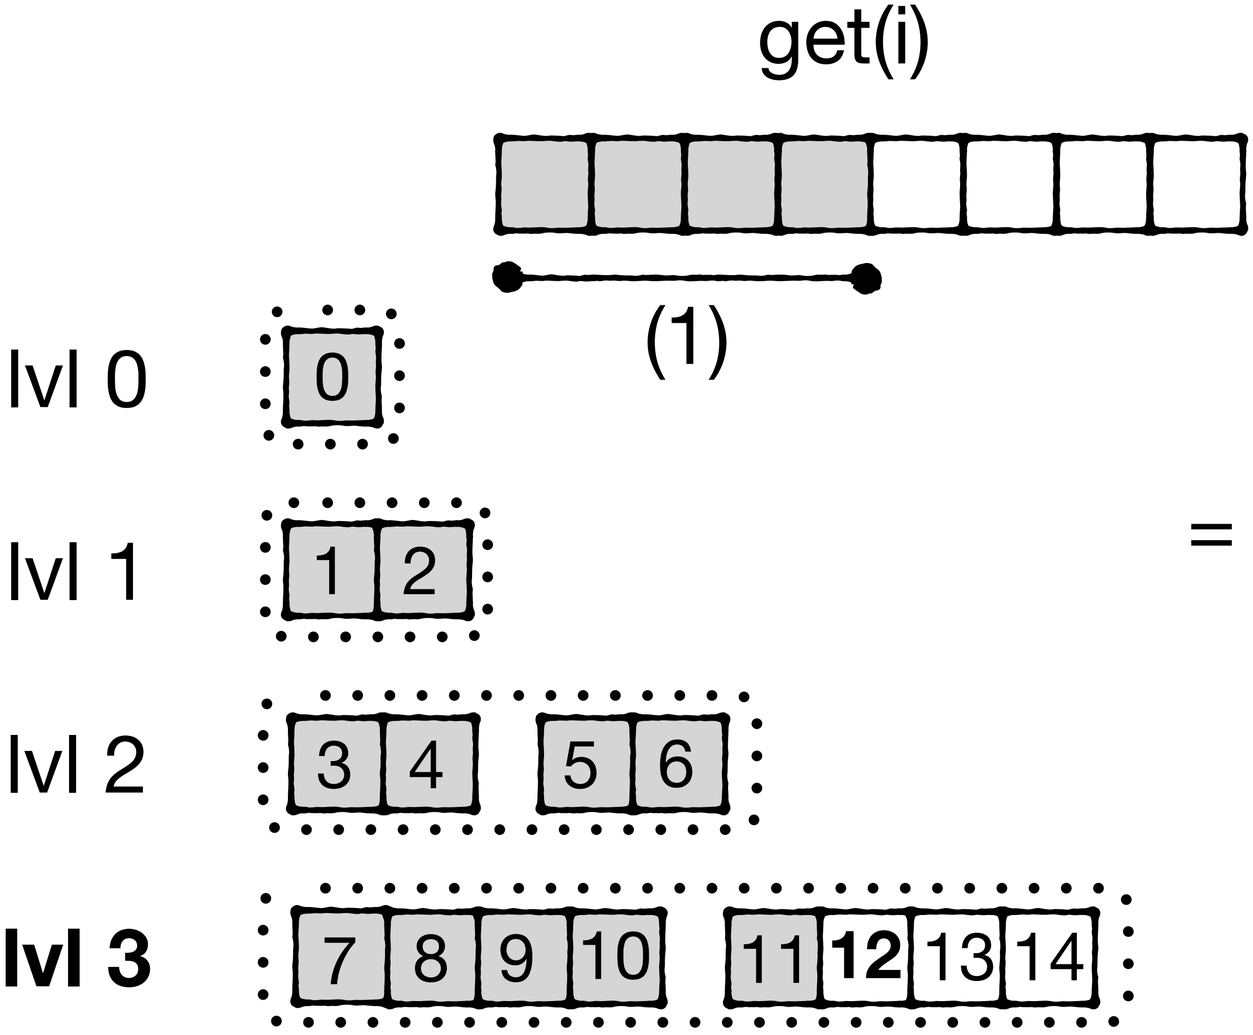

blocks in levels 0 to k-1

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

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

≈0.6 ns rather than ≈10ns

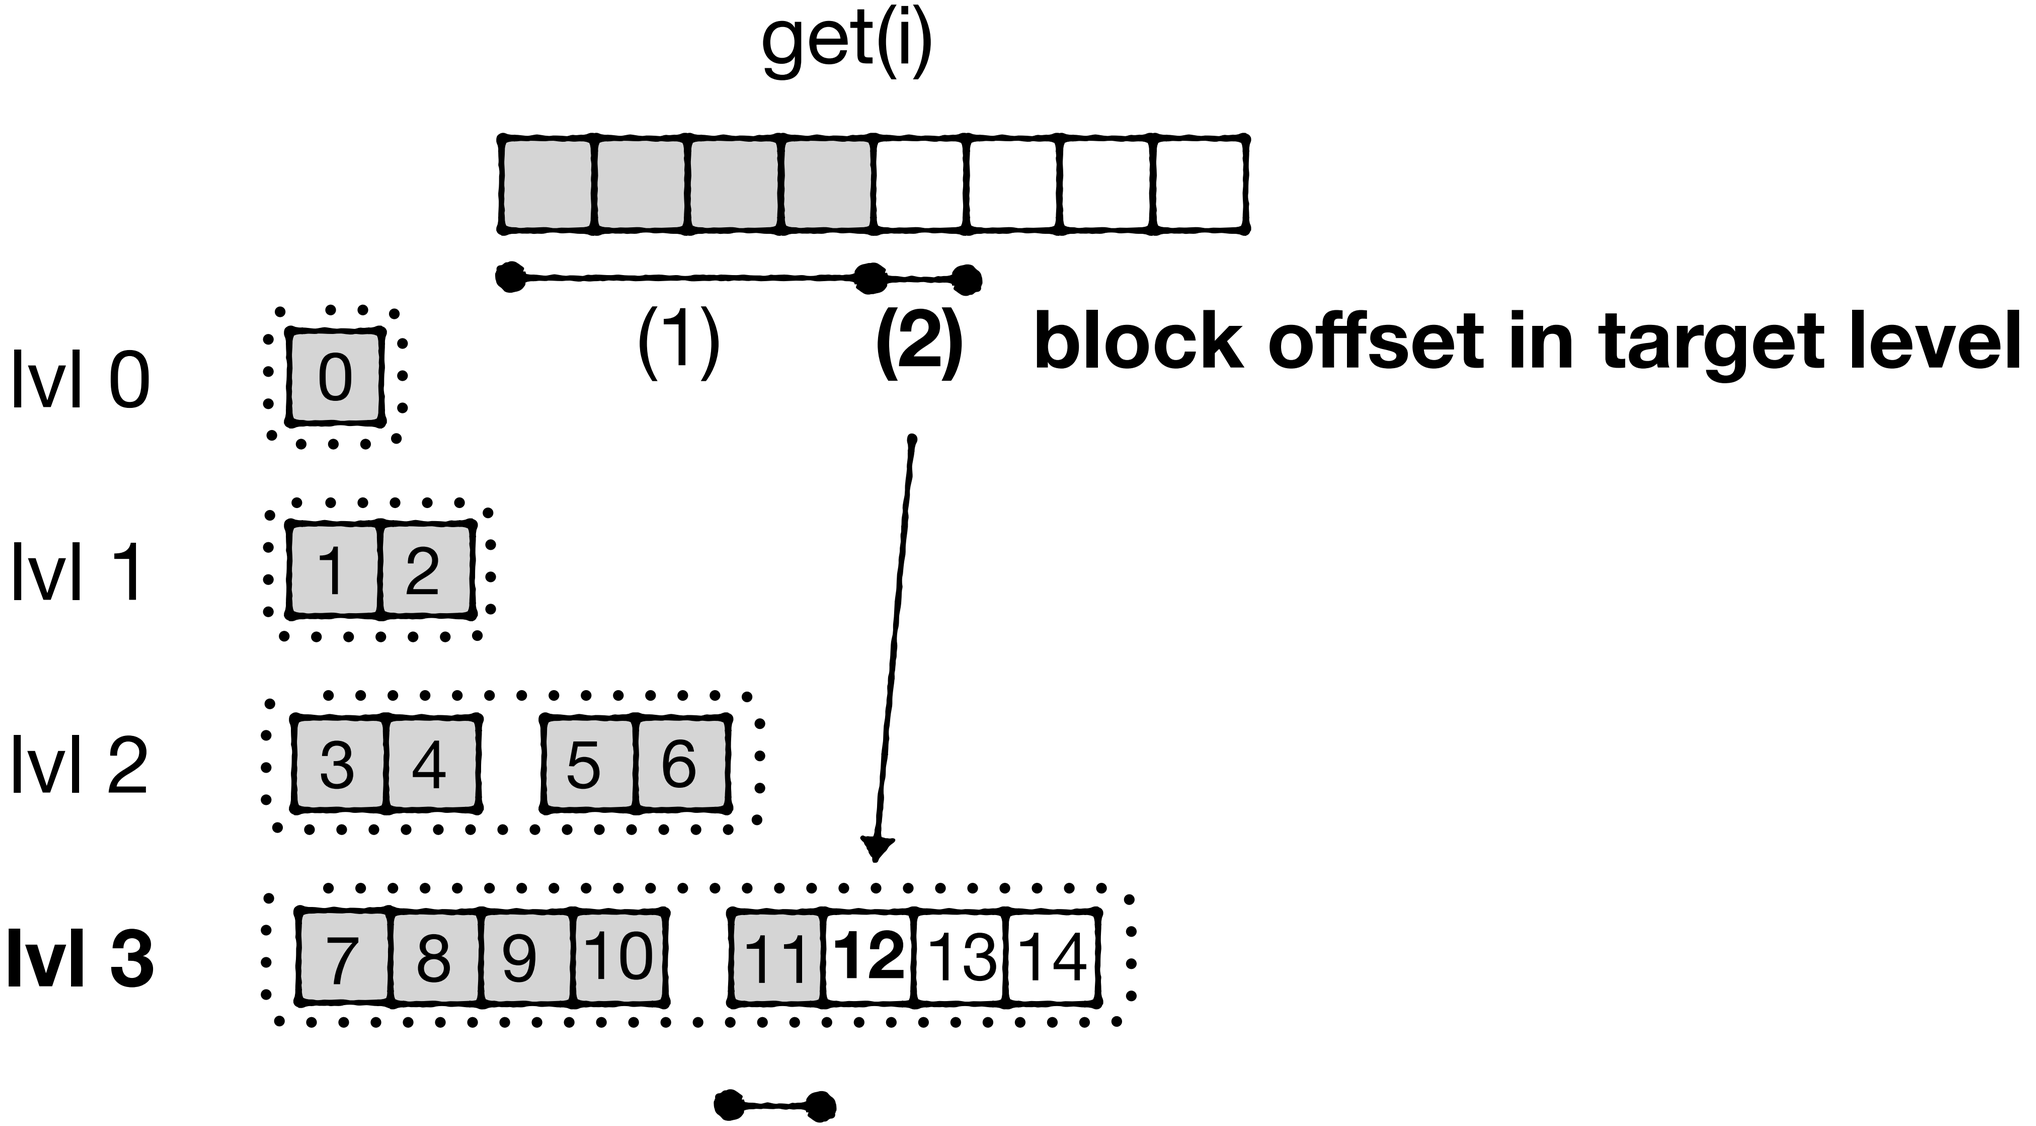
get(i)
(1)
(2)
block offset in target level
lvl 0
0
lvl 1
1
2
lvl 2
3
4
5
6
lvl 3
7
8
9
10
11
12
13
14
(3)
slot offset in target block

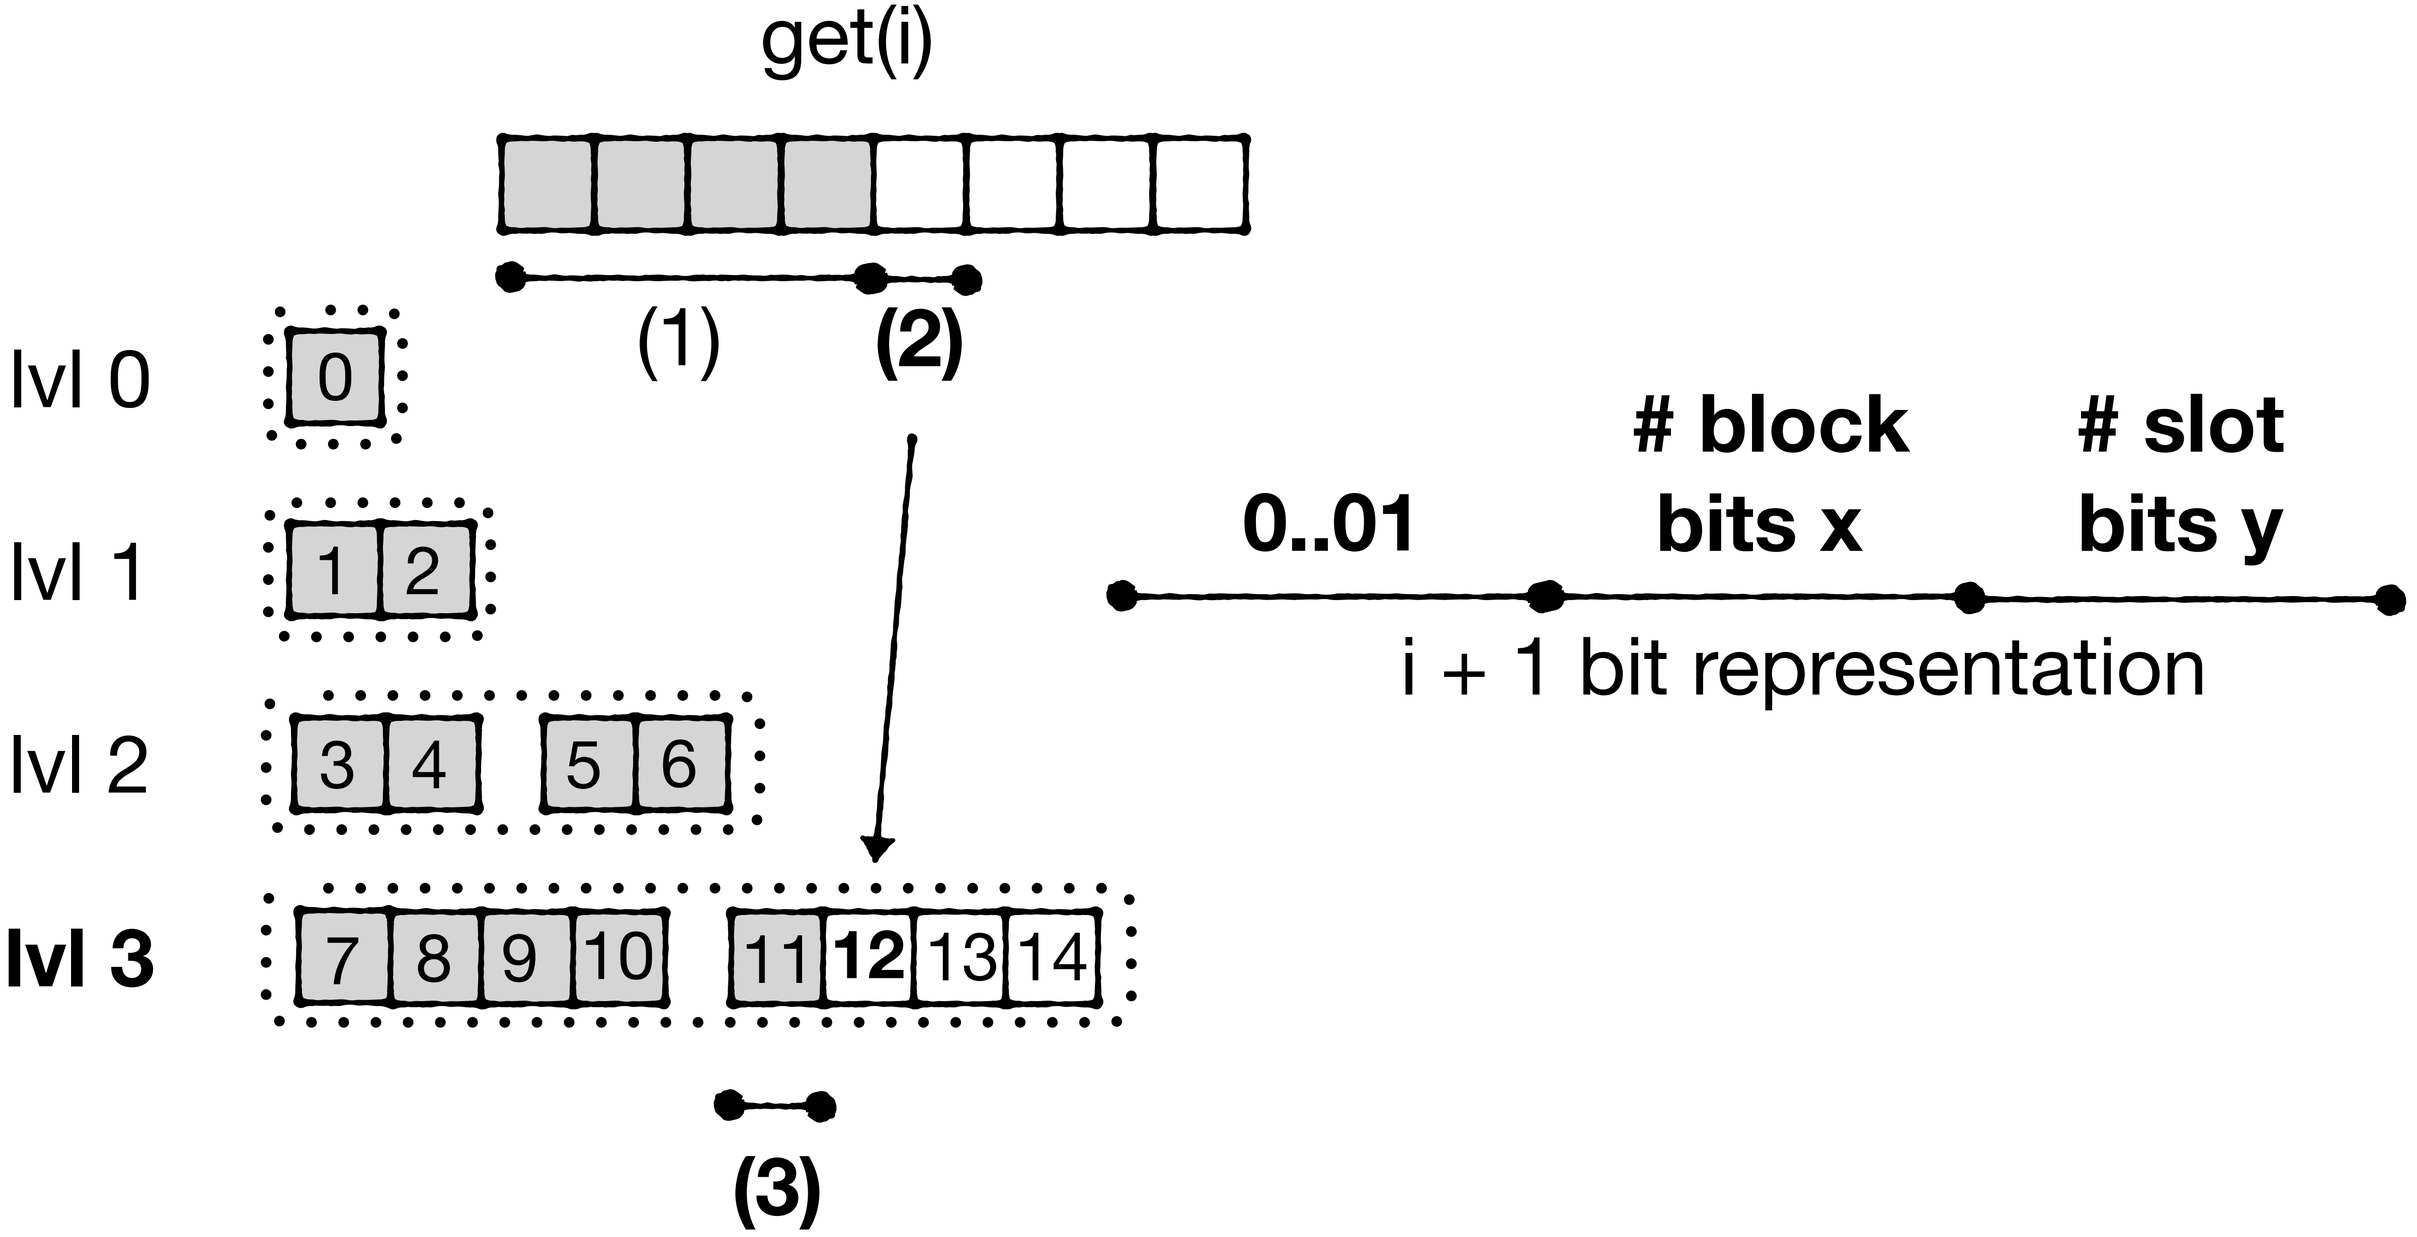
get(i)
(1)
(2)
lvl 0
0
lvl 1
1
2
lvl 2
3
4
5
6
lvl 3
7
8
9
10
11
12
13
14
(3)
0..01
block bits x
slot bits y
i + 1 bit representation

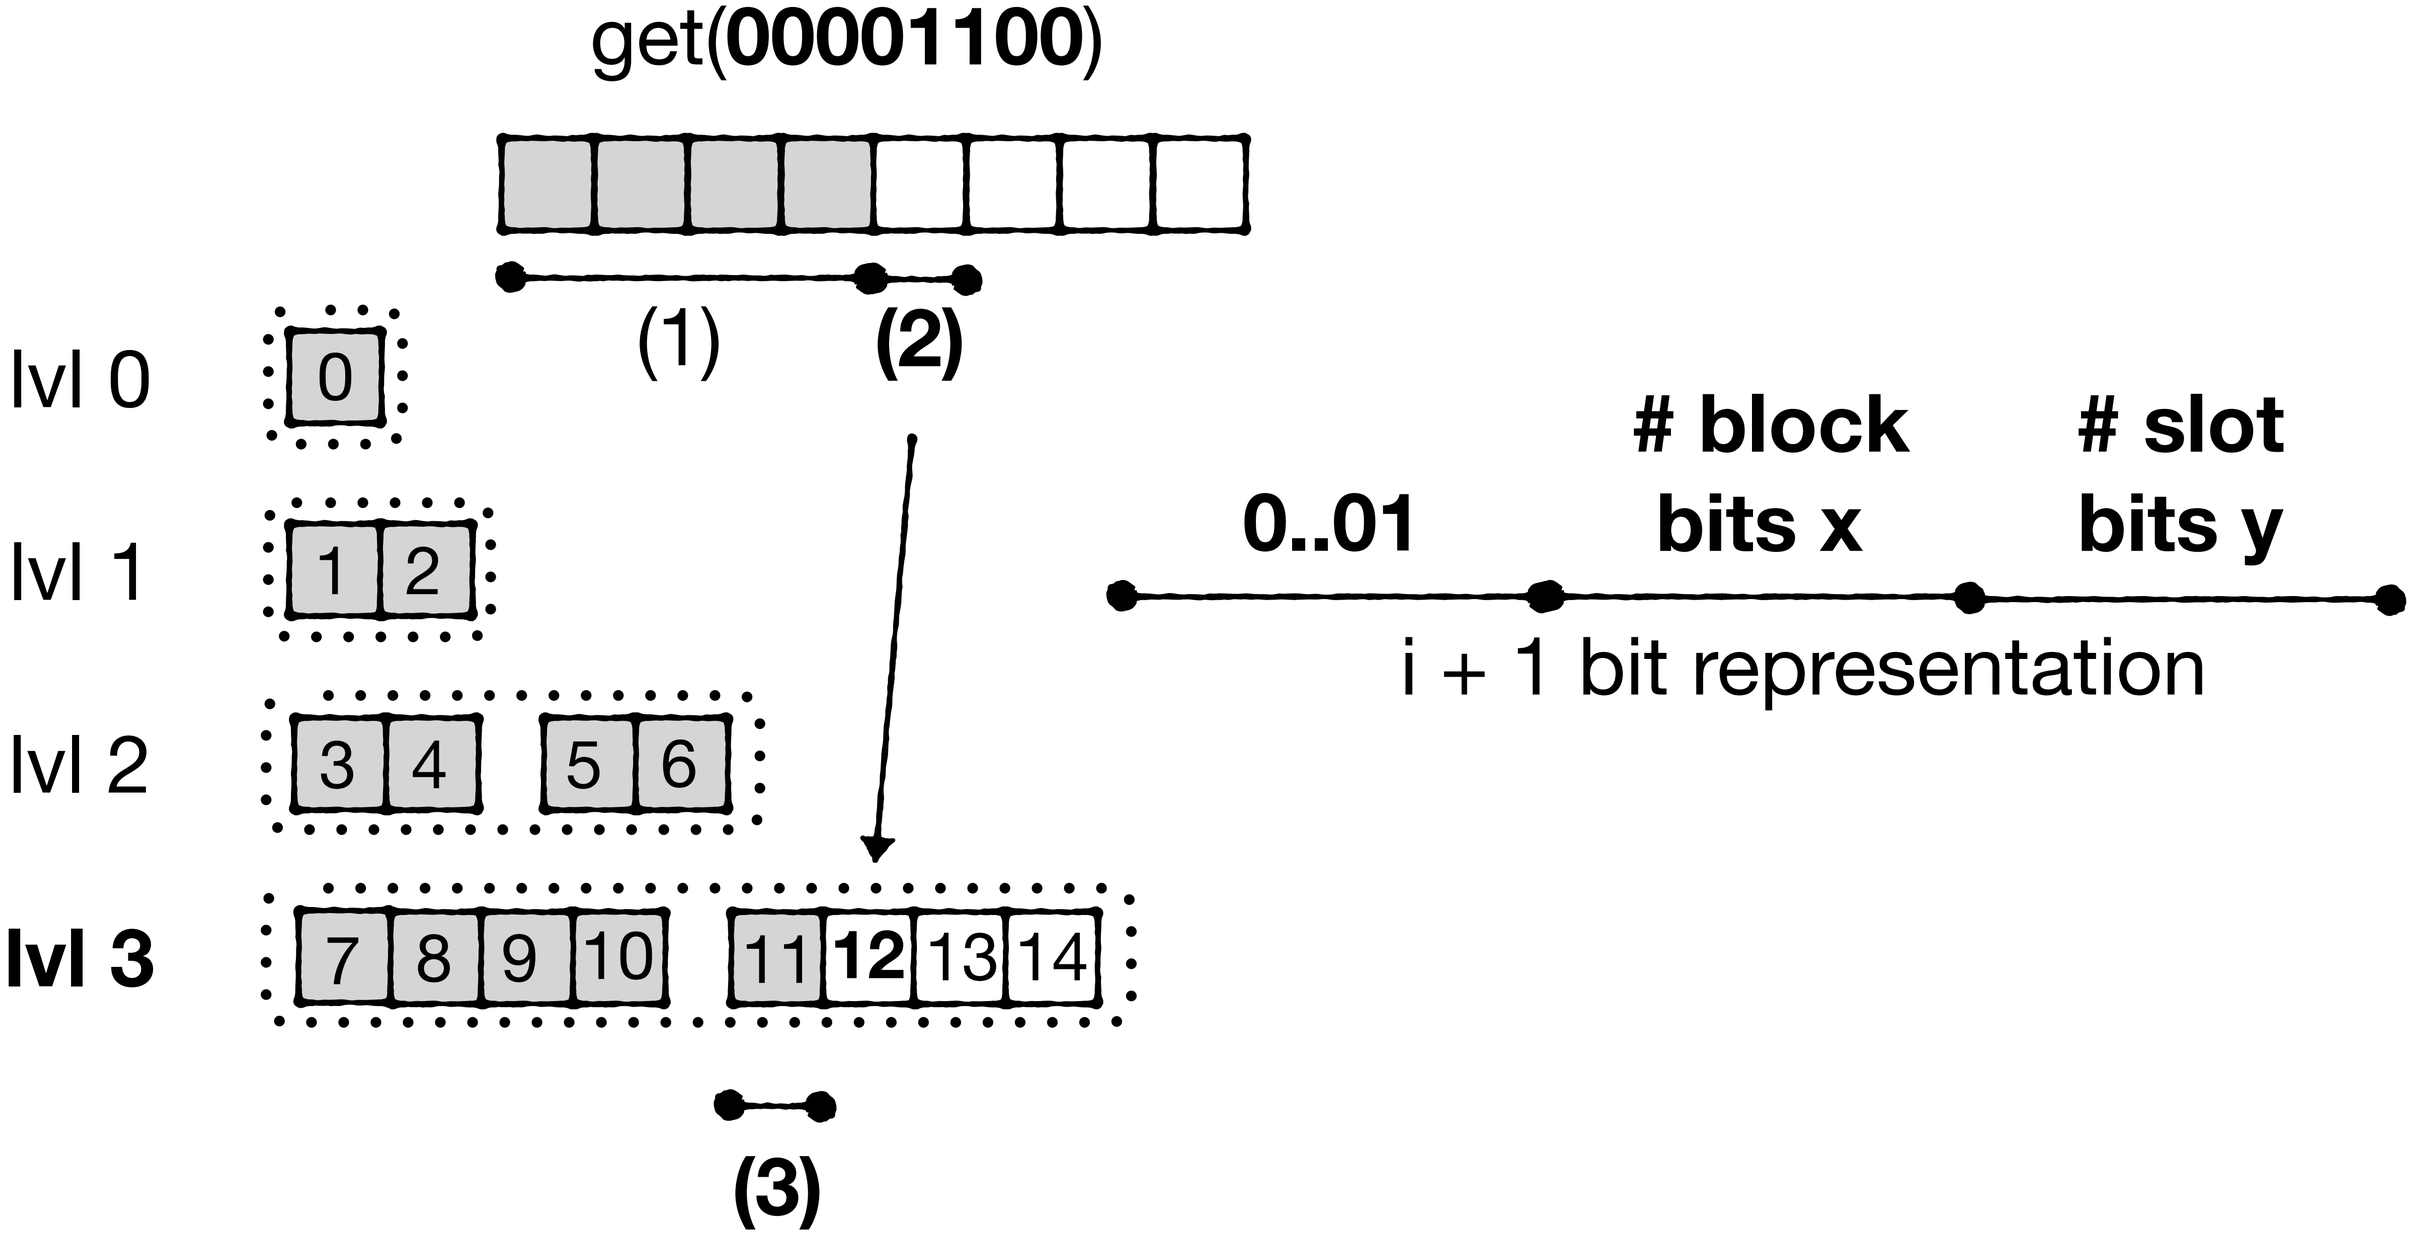
get(**00001100**)
(1)
(2)
lvl 0
0
lvl 1
1
2
lvl 2
3
4
5
6
lvl 3
7
8
9
10
11
12
13
14
(3)
 # block
bits x
 # slot
bits y
0..01
i + 1 bit representation

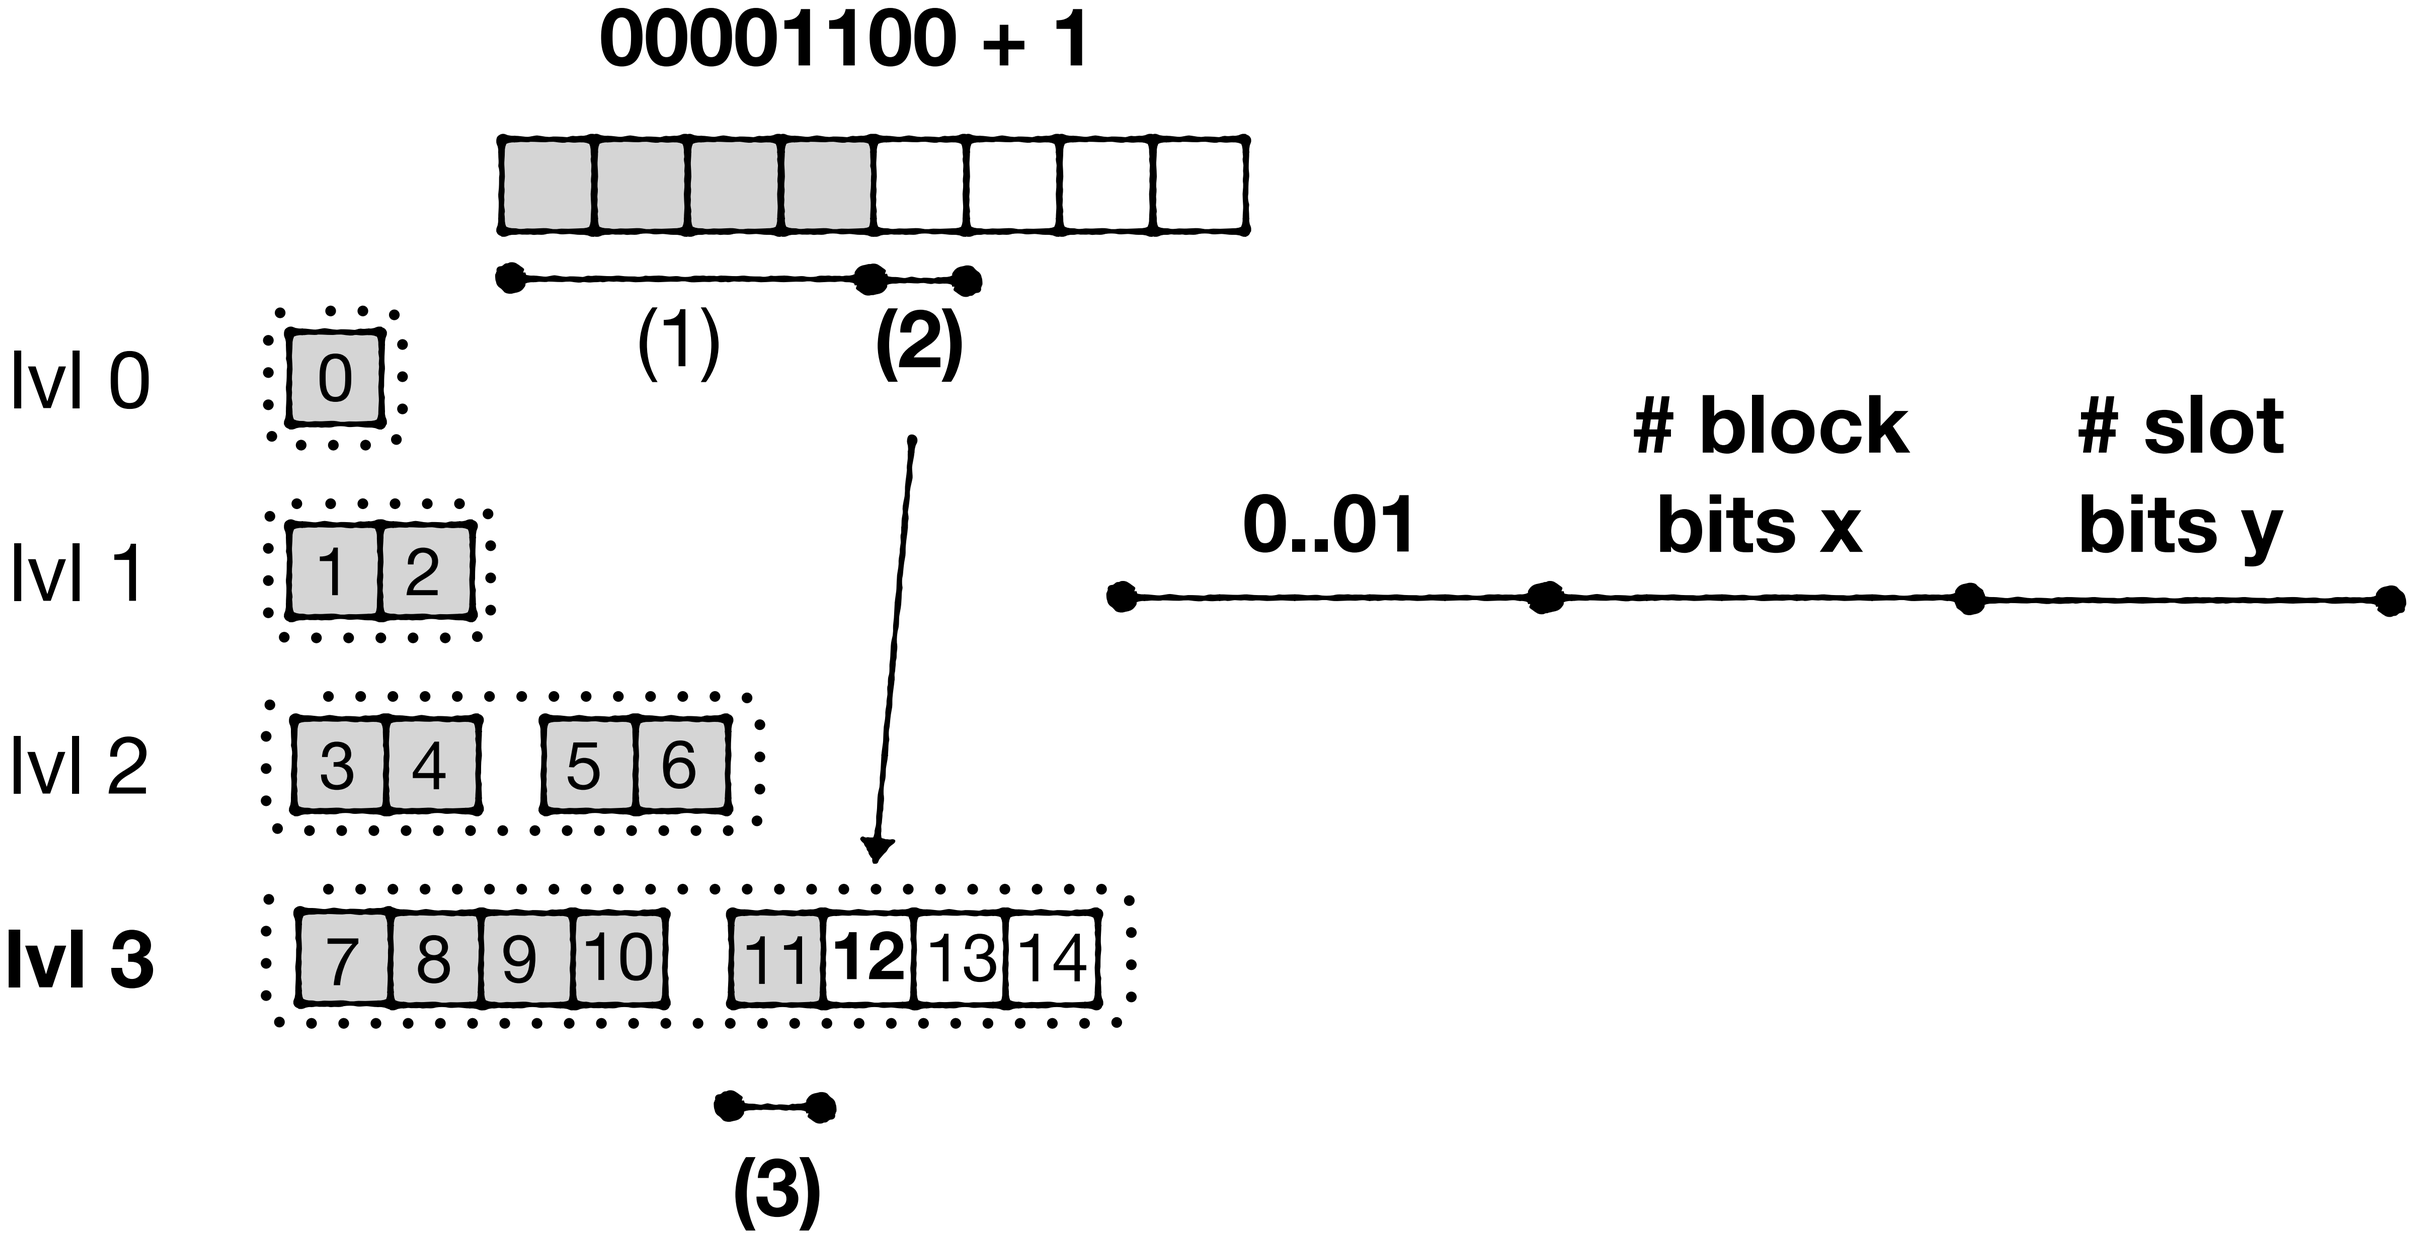
00001100 + 1
(1)
(2)
lvl 0
0
lvl 1
1 2
lvl 2
3 4 5 6
lvl 3
7 8 9 10 11 12 13 14
(3)
0..01
block bits x
slot bits y

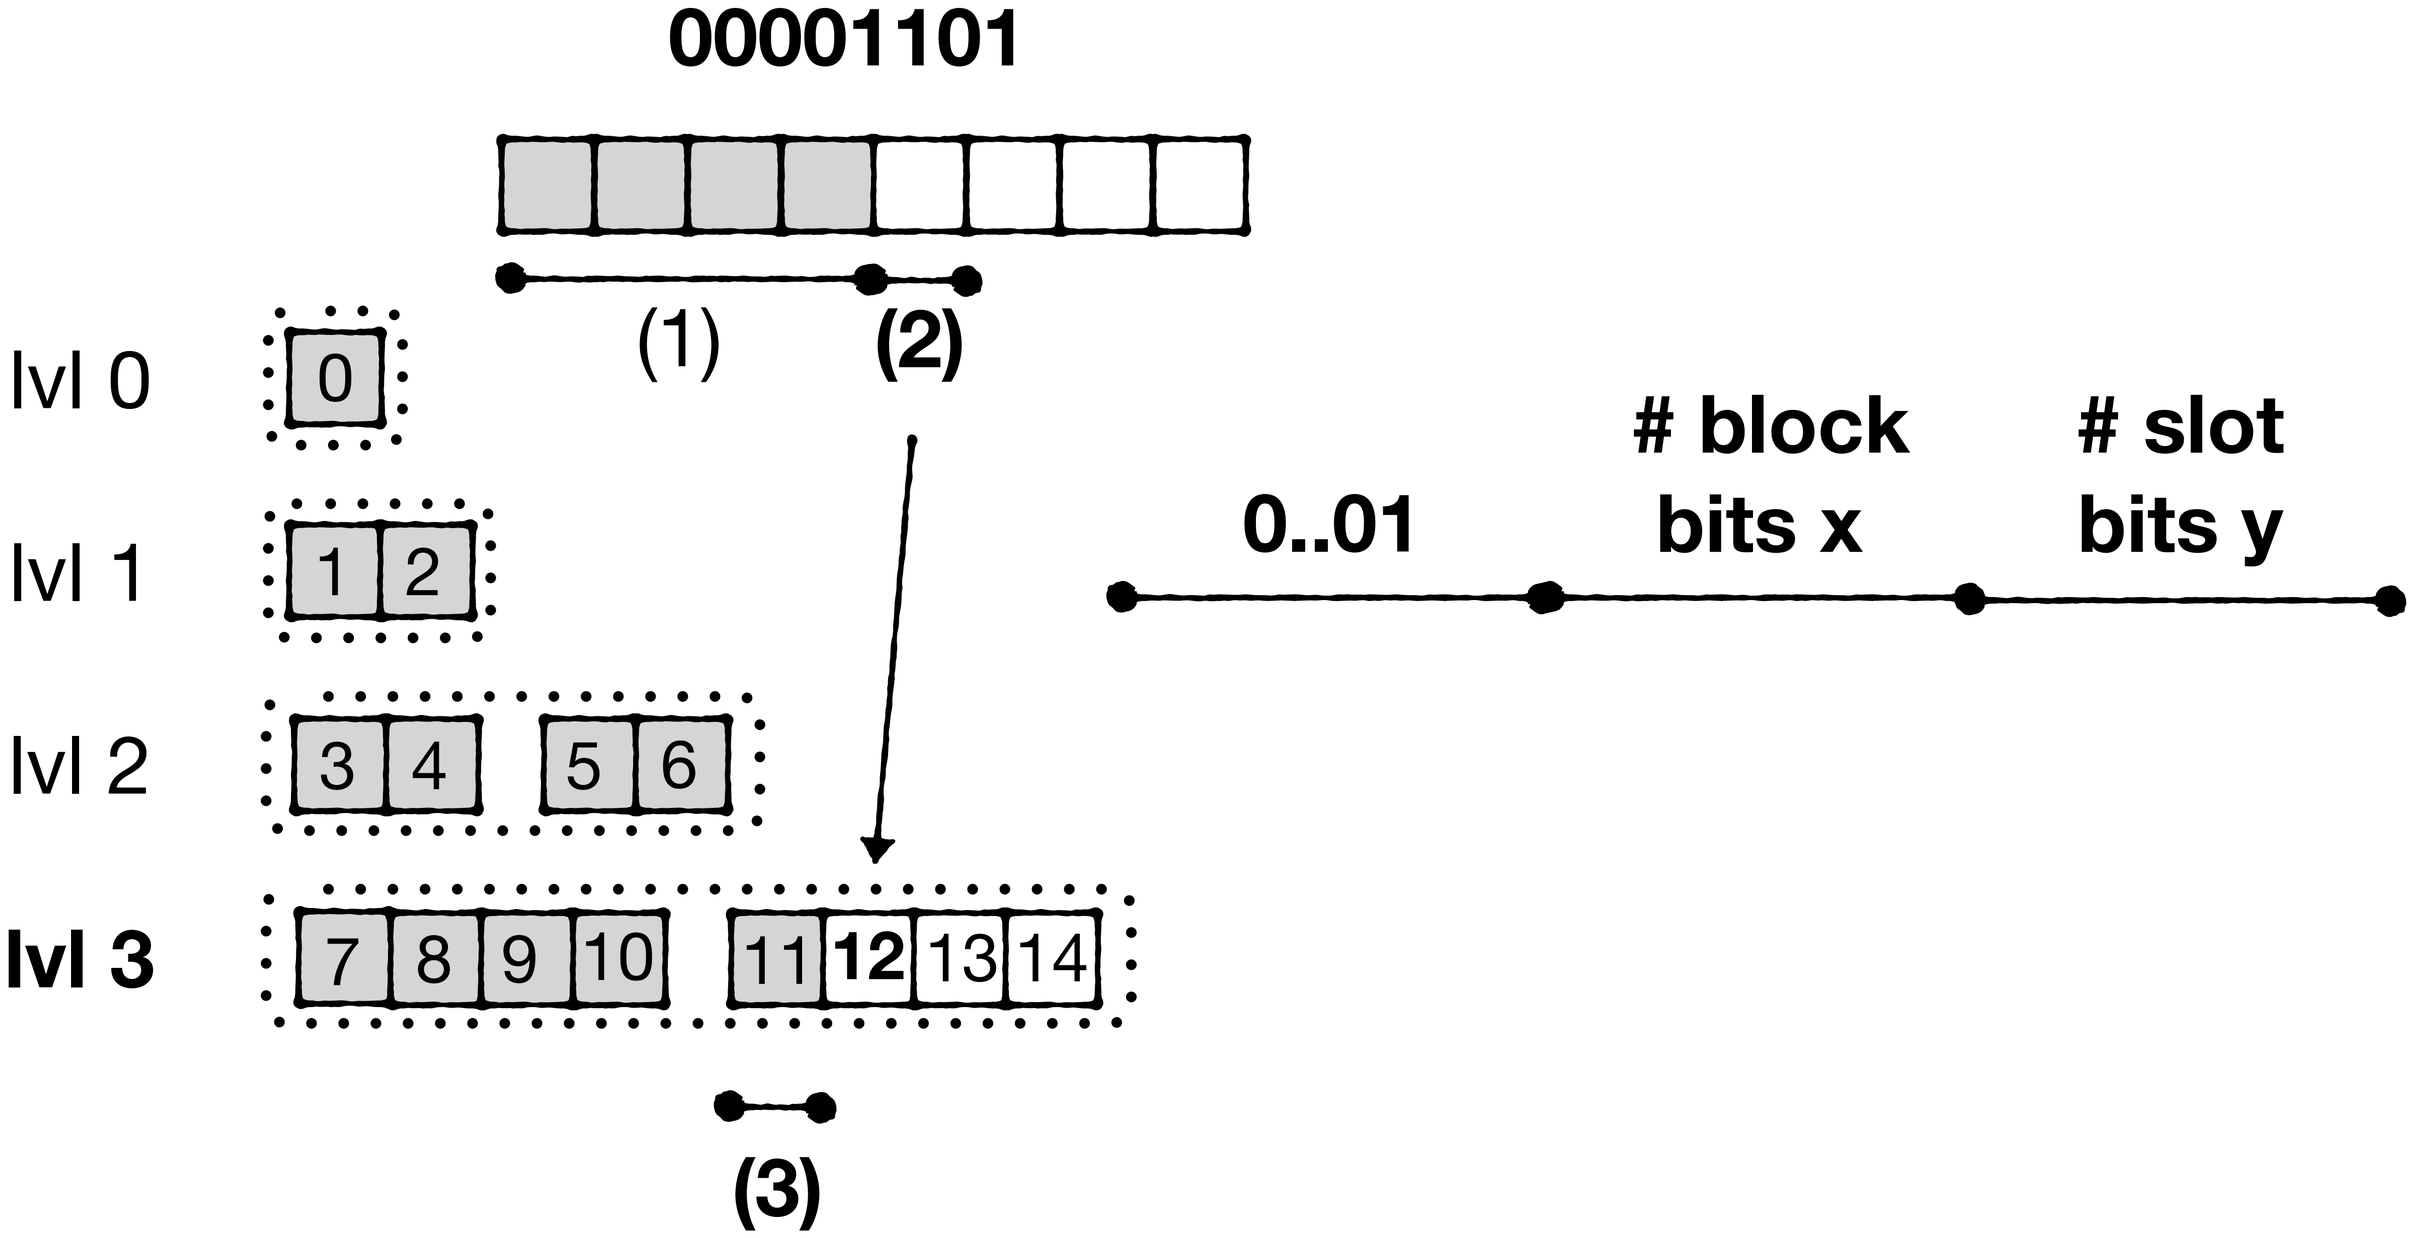
00001101
(1)
(2)
lvl 0
0
lvl 1
1
2
lvl 2
3
4
5
6
lvl 3
7
8
9
10
11
12
13
14
(3)
0..01
block bits x
slot bits y

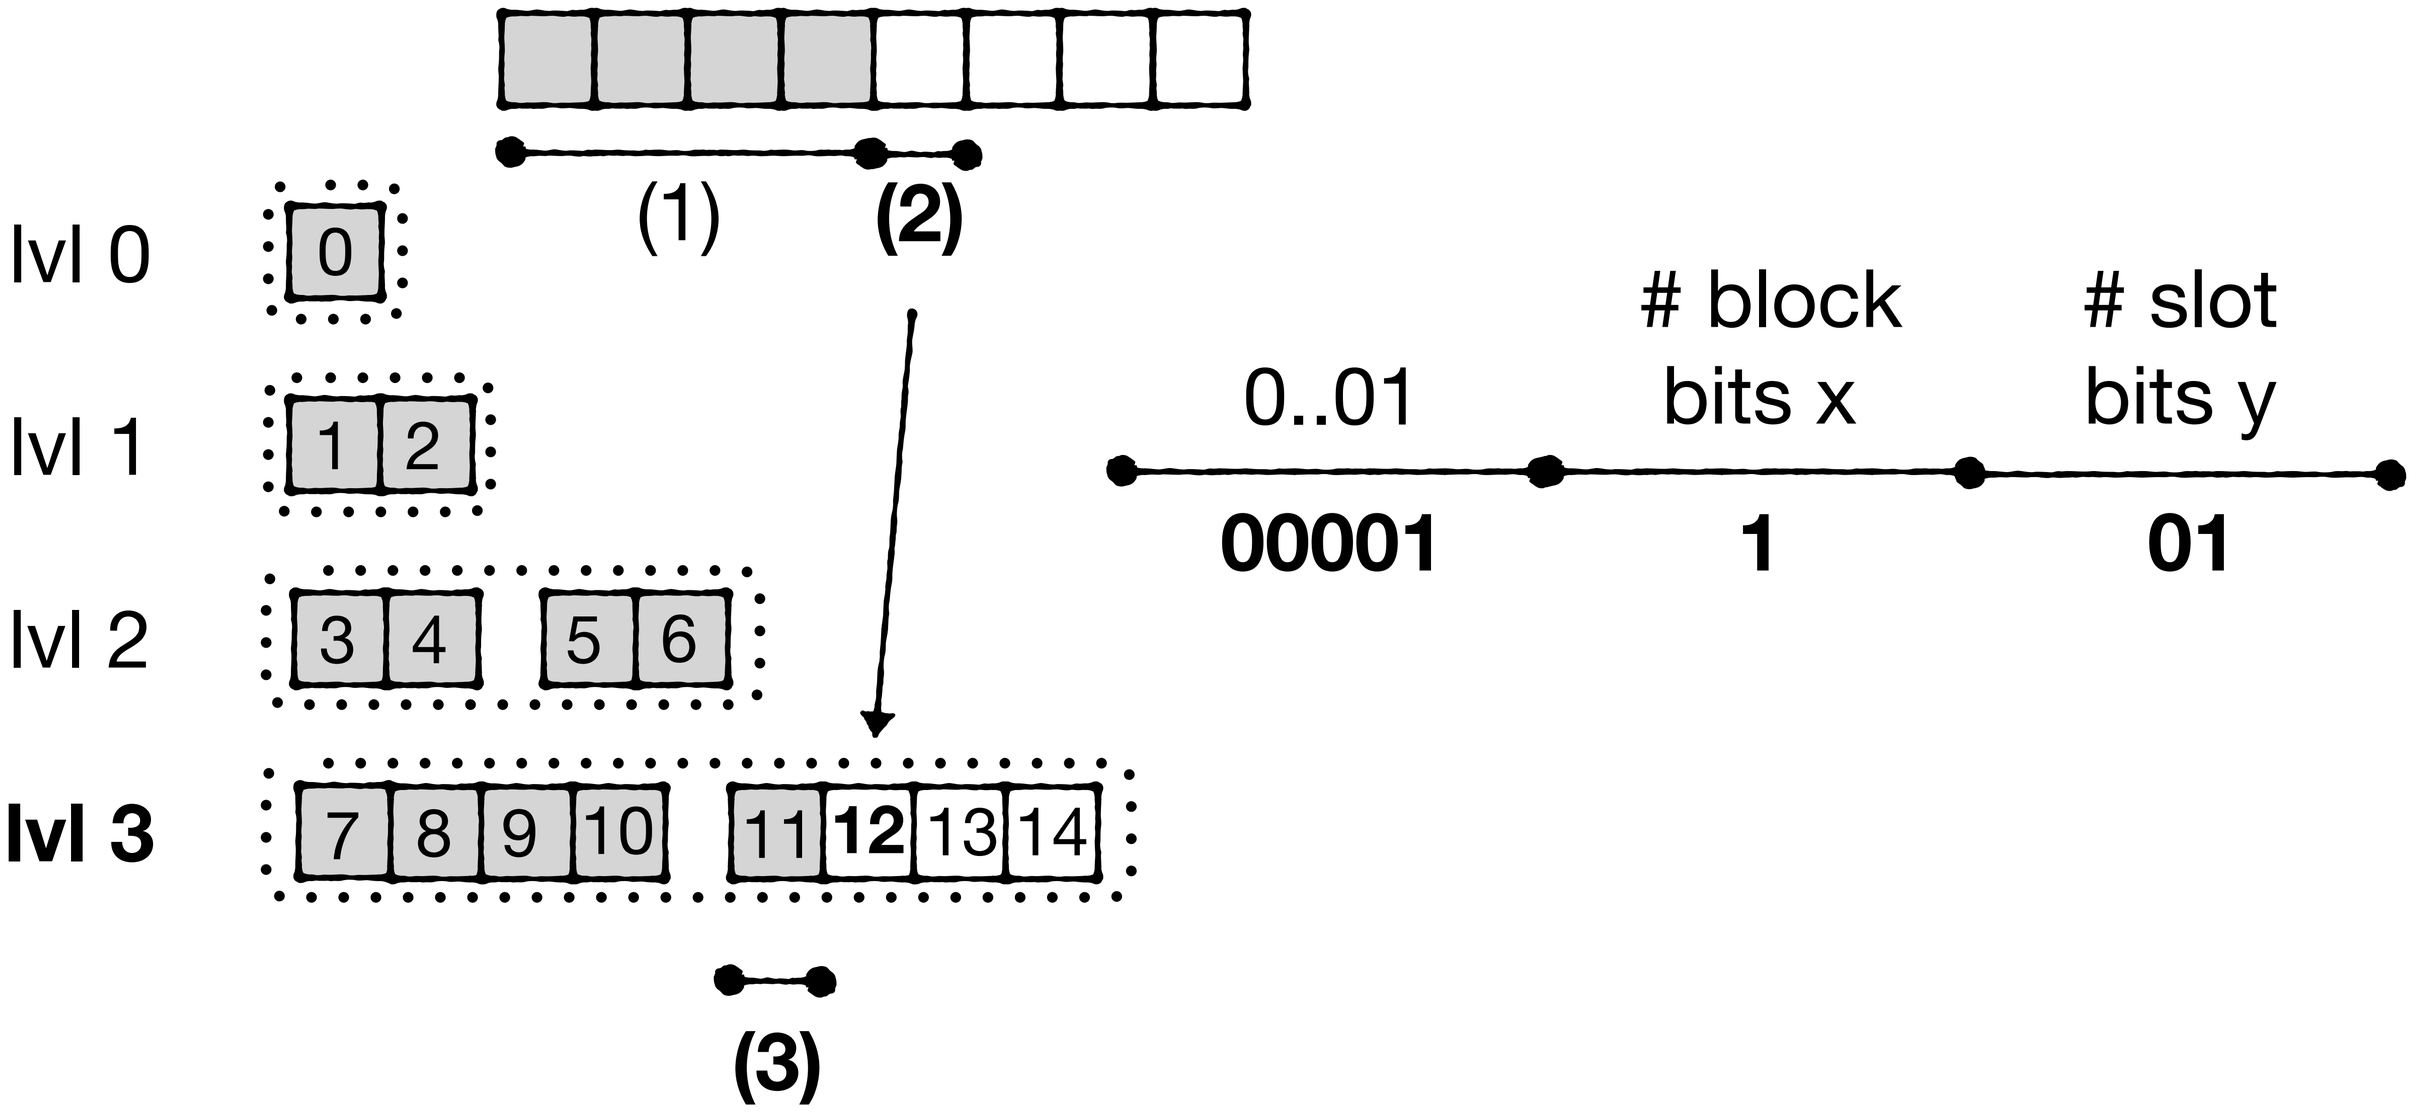
lvl 0
lvl 1
lvl 2
lvl 3
0
1 2
3 4 5 6
7 8 9 10 11 12 13 14
(1)
(2)
(3)
0..01
block bits x
slot bits y
00001
1
01

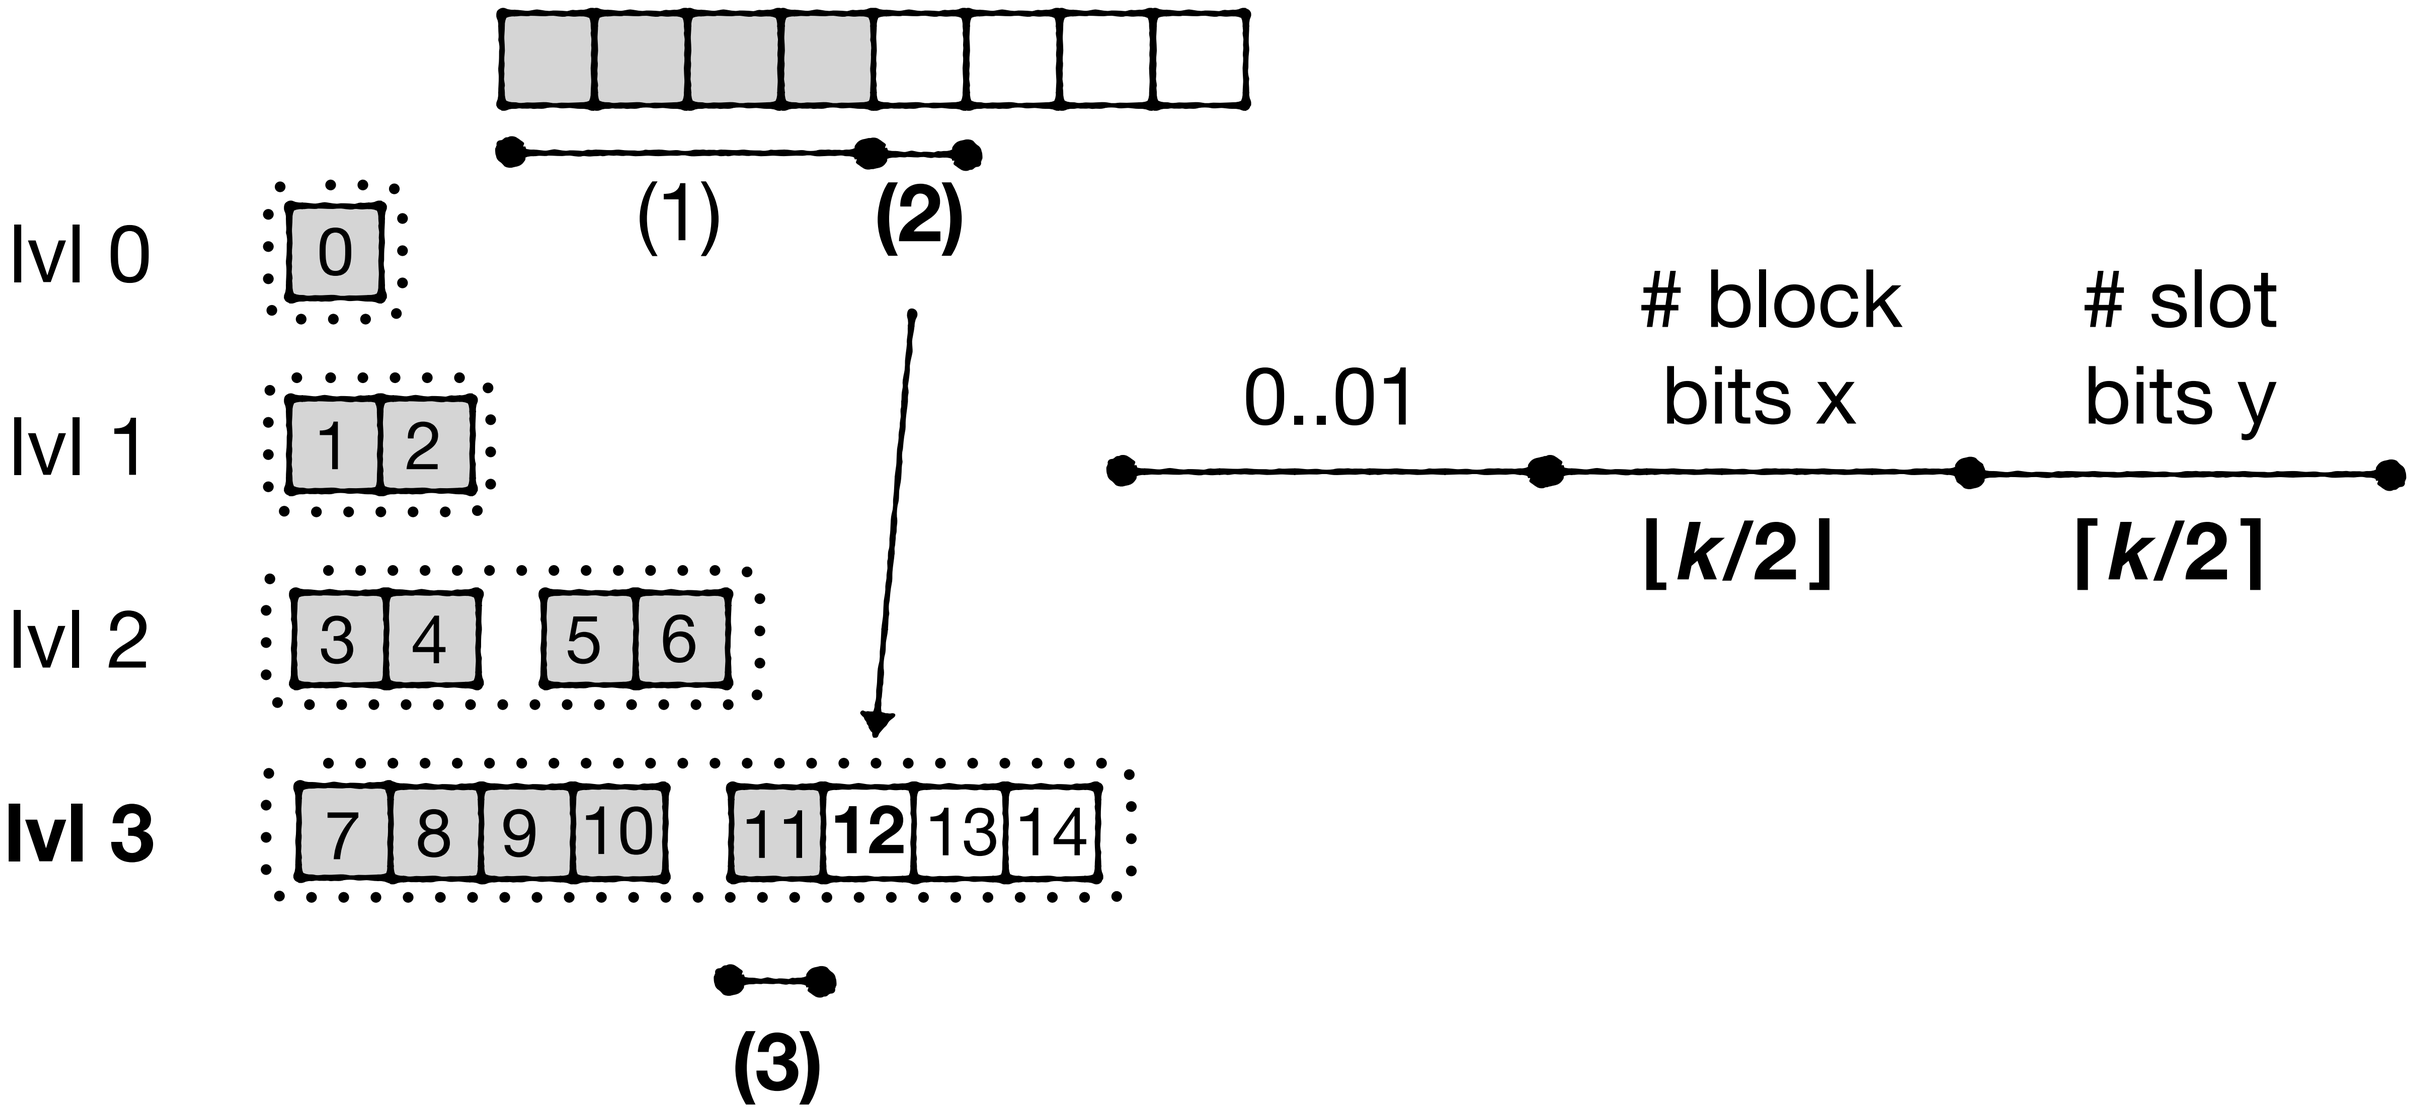
lvl 0
lvl 1
lvl 2
lvl 3
0
1 2
3 4 5 6
7 8 9 10 11 12 13 14
(1)
(2)
(3)
0..01
block bits x
slot bits y
⌊k/2⌋
⌈k/2⌉

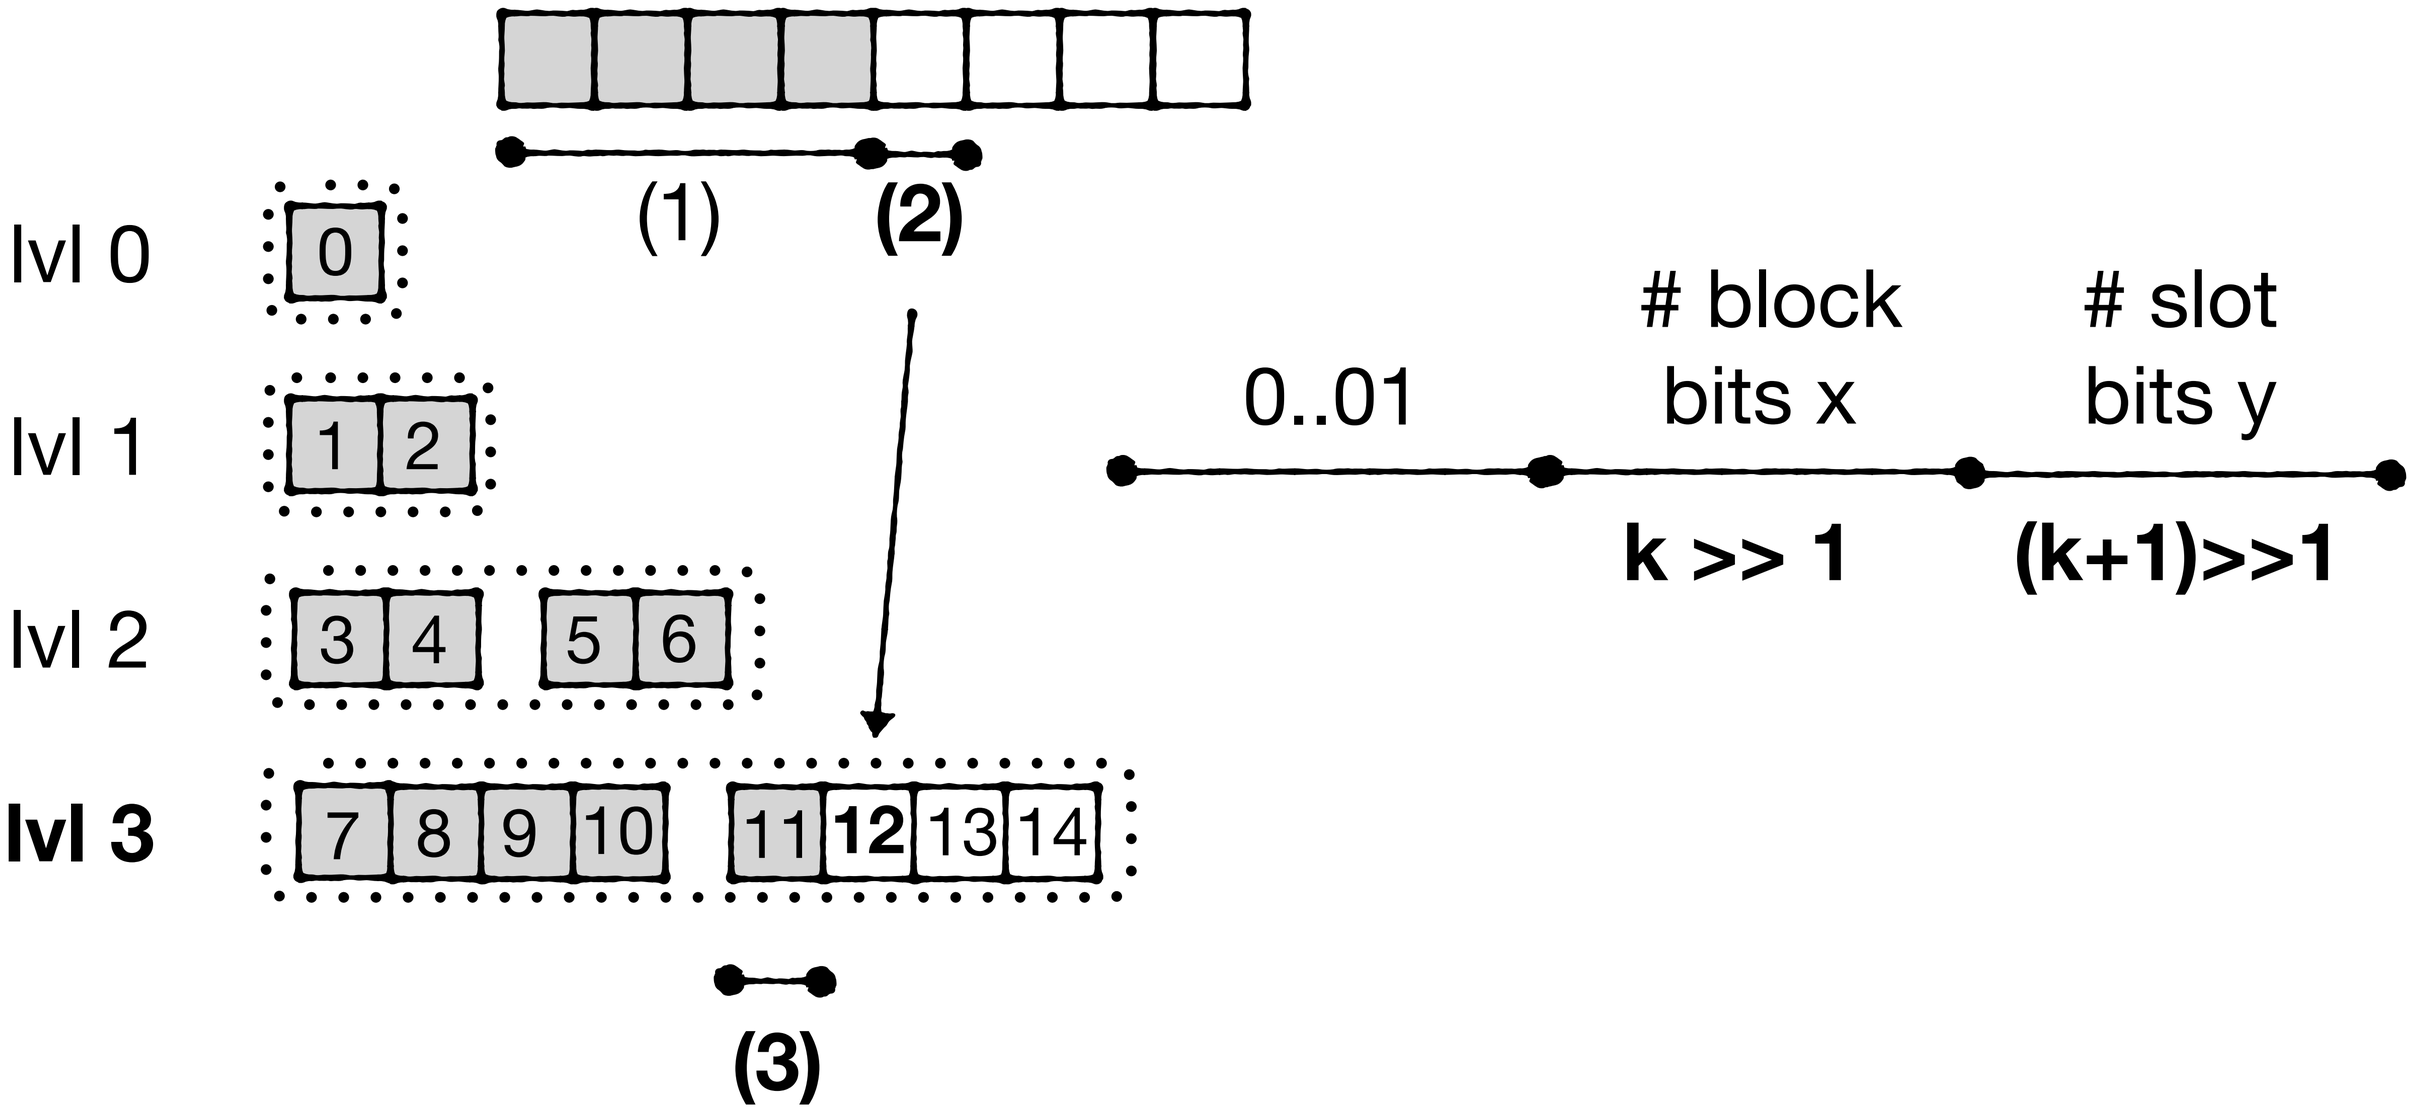
lvl 0
lvl 1
lvl 2
lvl 3
0
1 2
3 4 5 6
7 8 9 10 11 12 13 14
(1)
(2)
(3)
0..01
block bits x
slot bits y
k >> 1
(k+1)>>1

(1) **(2)**

lvl 0 0

lvl 1 1 2

lvl 2 3 4 5 6

lvl 3 7 8 9 10 11 **12** 13 14

(3)

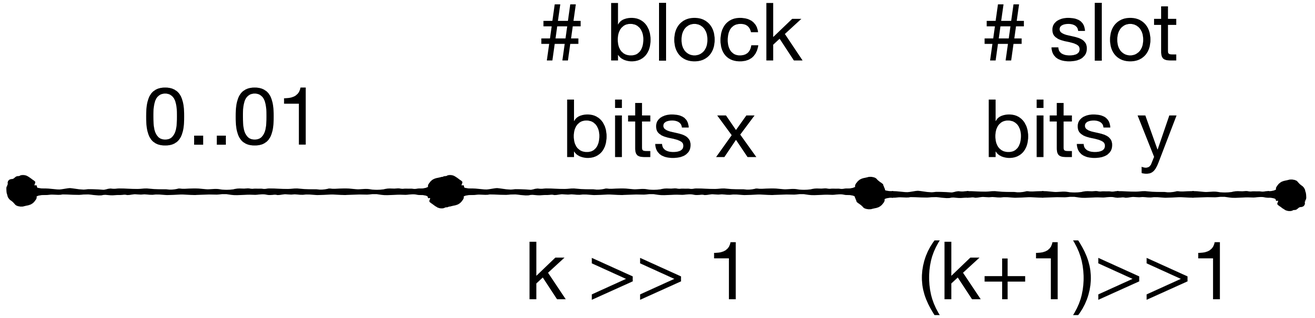

Slot offset **((i+1) & ((1 << y) - 1)**

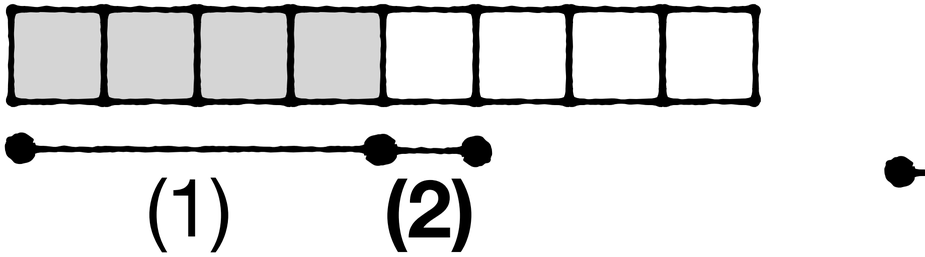

(1) **(2)**

lvl 0 0

lvl 1 1 2

lvl 2 3 4 5 6

lvl 3 7 8 9 10 11 **12** 13 14

(3)

0..01 # block bits x **# slot bits y**

k >> 1 (k+1)>>1

Slot offset ((i+1) & **((1 << y) - 1)**

Mask to only keep y least significant bits

(1) **(2)**

lvl 0 0

lvl 1 1 2

lvl 2 3 4 5 6

lvl 3 7 8 9 10 11 **12** 13 14

(3)

Slot offset ((i+1) & ((1 << y) - 1)

Block offset ((i+1) >> y) & ((1 << x) - 1)

(1) **(2)**

lvl 0 0

lvl 1 1 2

lvl 2 3 4 5 6

lvl 3 7 8 9 10 11 **12** 13 14

(3)

Slot offset ((i+1) & ((1 << y) - 1)

Block offset ((i+1) **>> y**) & ((1 << x) - 1)

Shift to least significant bits position

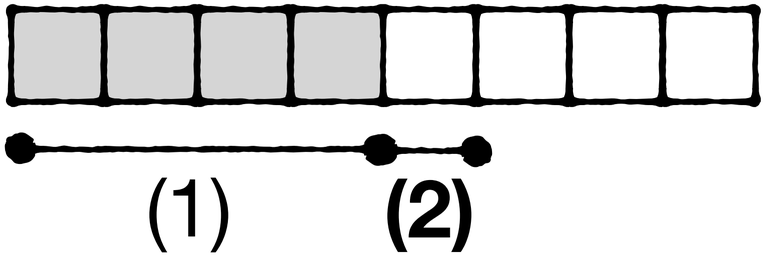

(1) **(2)**

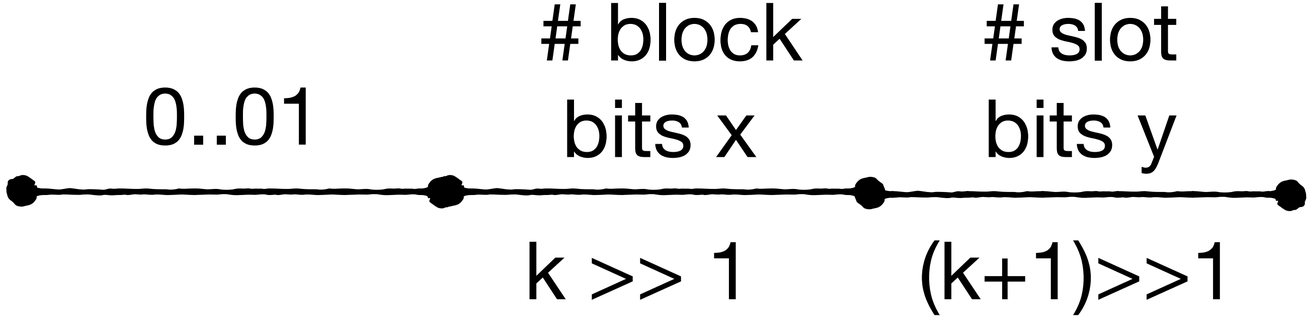

lvl 0 0

lvl 1 1 2

lvl 2 3 4 5 6

lvl 3 7 8 9 10 11 **12** 13 14

(3)

Slot offset ((i+1) & ((1 << y) - 1)

Block offset ((i+1) >> y) & **((1 << x) - 1)**

Mask to only keep x least significant bits

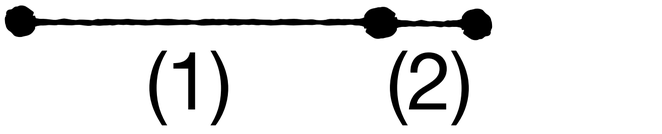

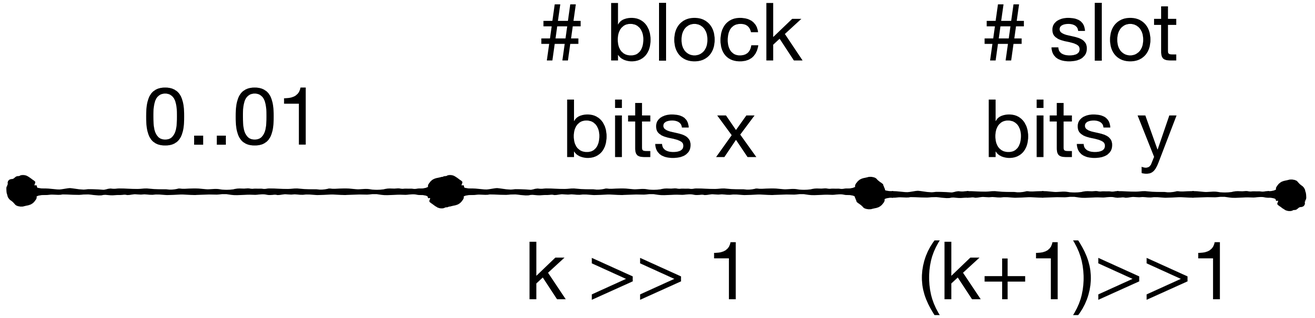

lvl 0

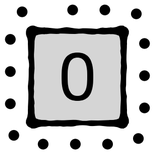

lvl 1

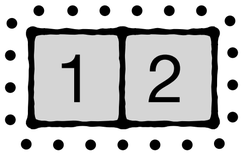

lvl 2

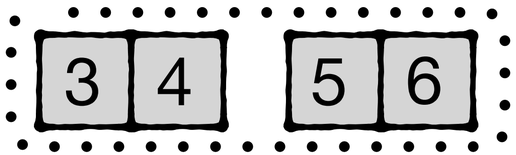

lvl 3

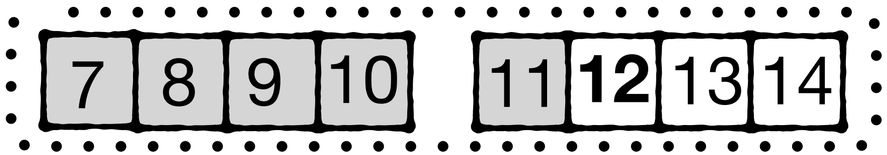

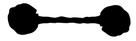

(3)

Slot offset ((i+1) & ((1 << y) - 1)

Block offset ((i+1) >> y) & ((1 << x) - 1)

We’re done :)

	Write-amp	Space-amp	Read-amp
indirection	≈1	$O(1+N^{-0.5})$	O(1)

	Write-amp	Space-amp	Read-amp
indirection	≈1	$O(1+N^{-0.5})$	O(1)
No indirection	$\frac{G}{G-1}$	O(G)	1

	Write-amp	**Space-amp**	**Read-amp**
indirection	≈1	$O(1+N^{-0.5})$	O(1)
No indirection **G > ϕ**	$\frac{\mathbf{G}}{\mathbf{G-1}}$	$\frac{\mathbf{G}}{\mathbf{G-1}} + \mathbf{G}$	1
No indirection **G < ϕ**	$\frac{\mathbf{G}}{\mathbf{G-1}}$	**G to G+1**	1

Thank you :)