



6.172
Performance
Engineering of
Software Systems

LECTURE 10
**Dynamic Storage
Allocation**

Charles E. Leiserson

October 12, 2010

Stack Allocation

Array and pointer

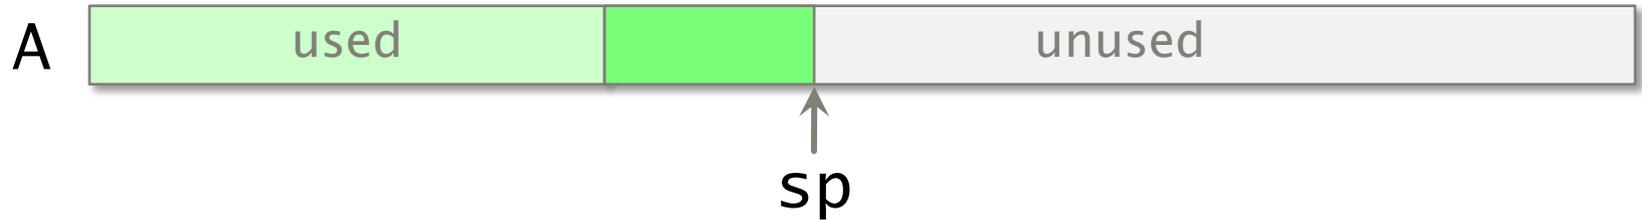


Allocate x bytes

```
sp += x;  
return A[sp - x];
```

Stack Allocation

Array and pointer



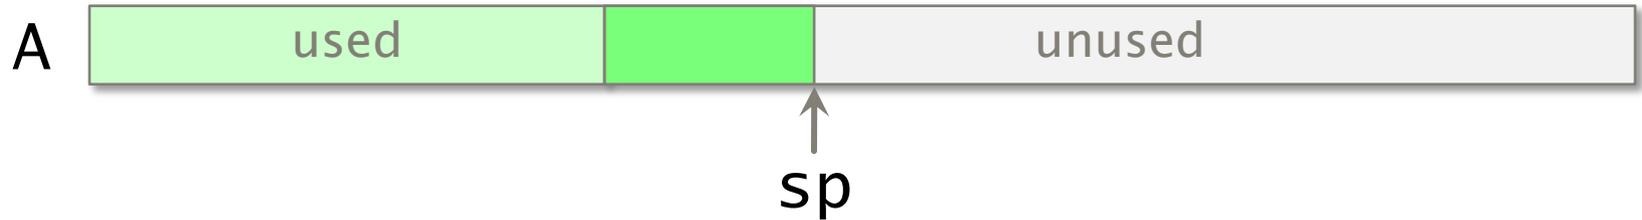
Allocate x bytes

```
sp += x;  
return A[sp - x];
```

Should check for
stack overflow.

Stack Deallocation

Array and pointer



Allocate x bytes

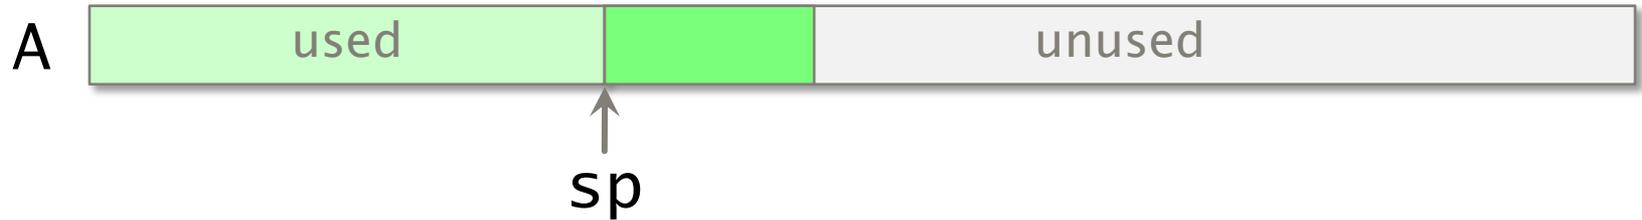
```
sp += x;  
return A[sp - x];
```

Free x bytes

```
sp -= x;
```

Stack Deallocation

Array and pointer



Allocate x bytes

```
sp += x;  
return A[sp - x];
```

Free x bytes

```
sp -= x;
```

Should check for
stack underflow.

Stack Storage

Array and pointer



Allocate x bytes

```
sp += x;  
return A[sp - x];
```

Free x bytes

```
sp -= x;
```

- Allocating and freeing take $\Theta(1)$ time.
- Must free consistent with stack discipline.
- Limited applicability, but great when it works!
- One can allocate on the call stack using `alloca()`, but this function is deprecated, and the compiler is more efficient with fixed-size frames.

Heap Allocation

- C provides `malloc()` and `free()`.
- C++ provides `new` and `delete`.

Unlike Java and Python, C and C++ provide no *garbage collector*. Heap storage allocated by the programmer must be freed explicitly.

Failure to do so creates a *memory leak*. Also, watch for *dangling pointers* and *double freeing*.

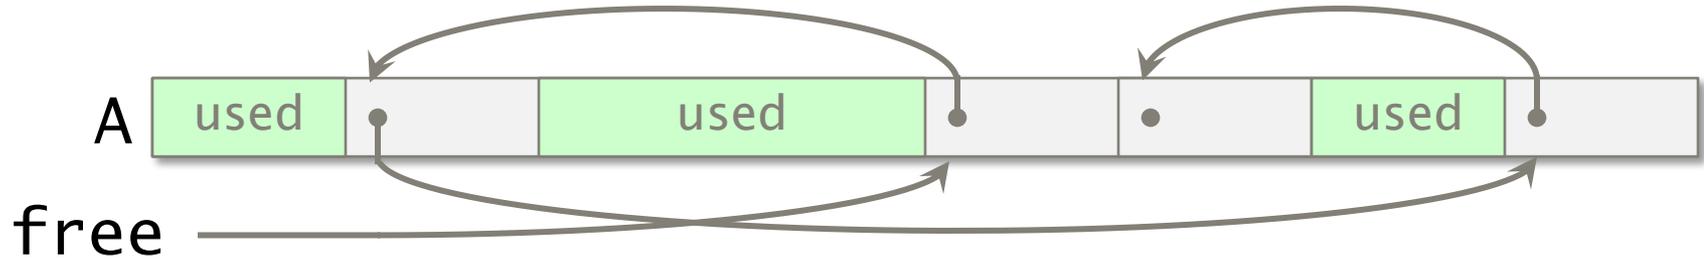
Memory checkers can assist in finding these pernicious bugs:

```
% valgrind --leakcheck=yes ./myprog <arguments>
```

Valgrind is installed on cloud machines. See <http://valgrind.org> for details.

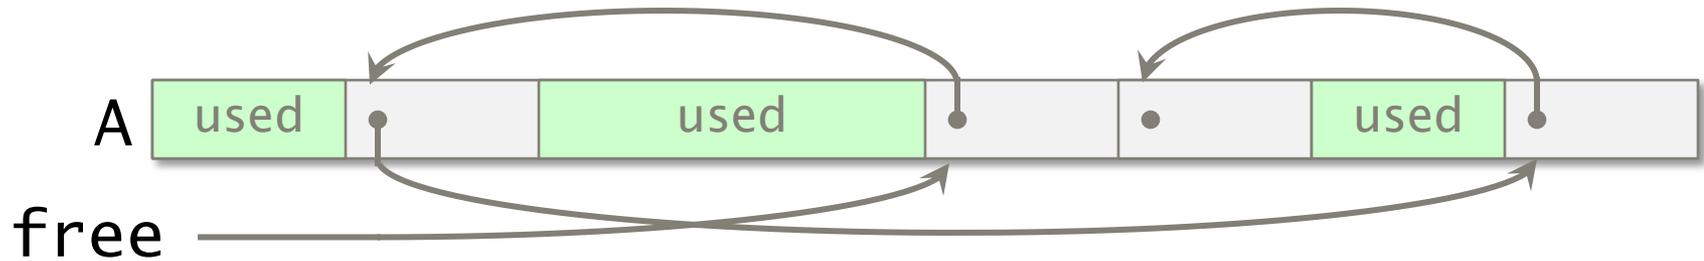
Fixed-Size Allocation

Free list



Fixed-Size Allocation

Free list

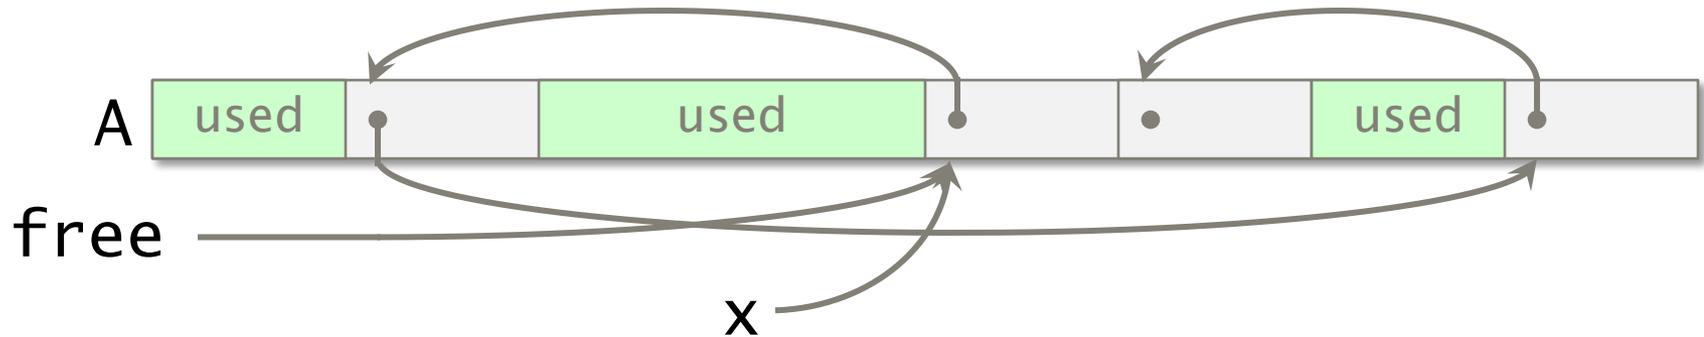


Allocate 1 object

```
x = free;  
free = free->next;  
return x;
```

Fixed-Size Allocation

Free list

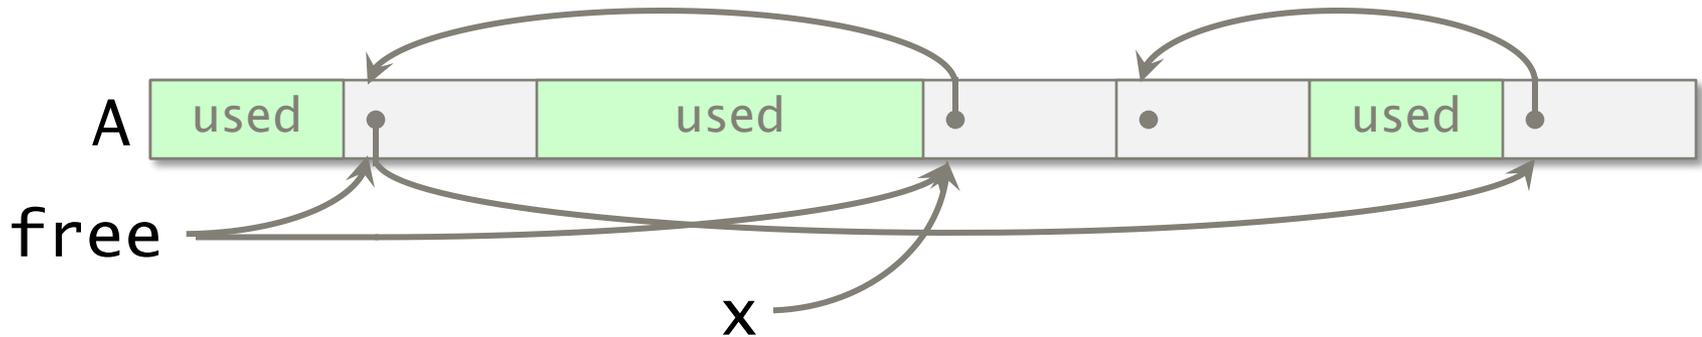


Allocate 1 object

```
x = free;  
free = free->next;  
return x;
```

Fixed-Size Allocation

Free list



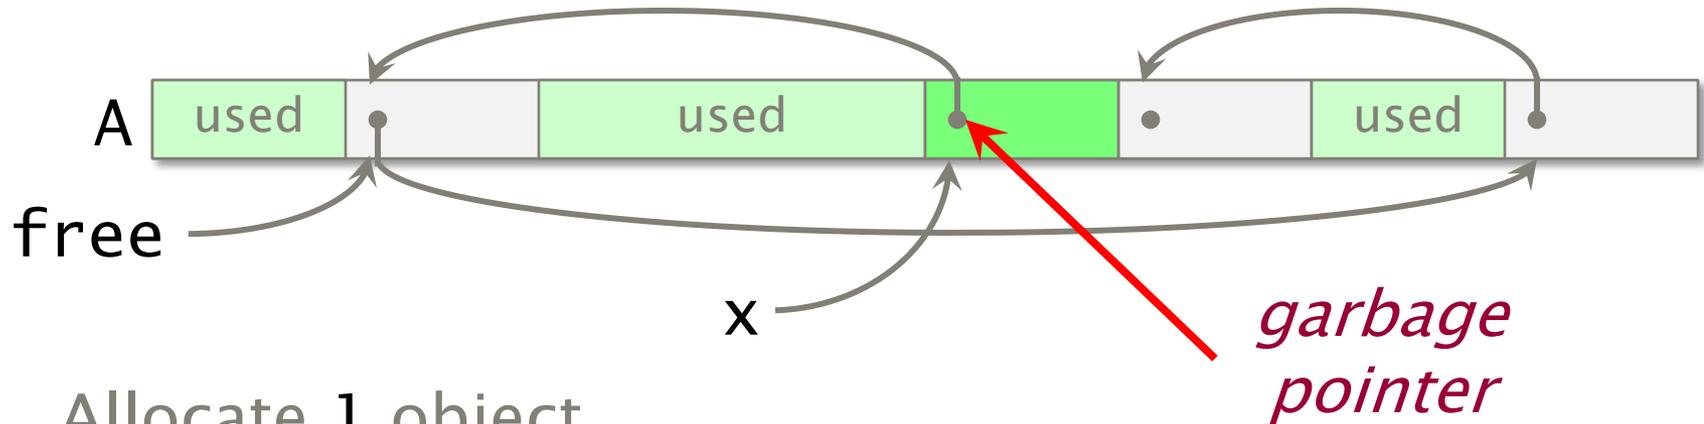
Allocate 1 object

```
x = free;  
free = free->next;  
return x;
```

Should check
`free != NULL.`

Fixed-Size Allocation

Free list

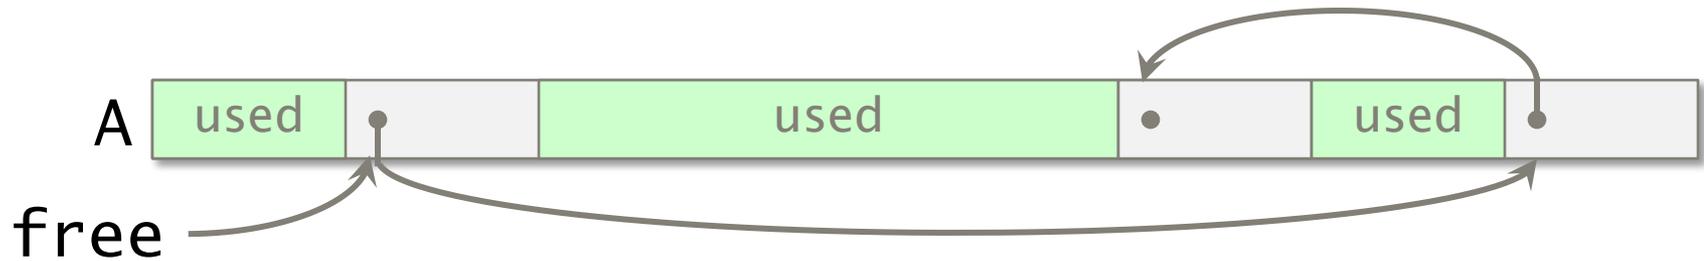


Allocate 1 object

```
x = free;  
free = free->next;  
return x;
```

Fixed-Size Allocation

Free list

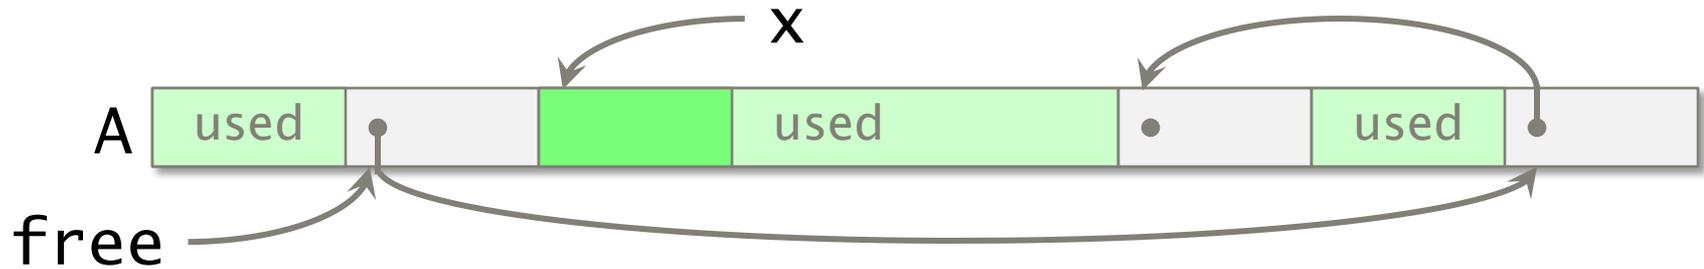


Allocate 1 object

```
x = free;  
free = free->next;  
return x;
```

Fixed-Size Deallocation

Free list



Allocate 1 object

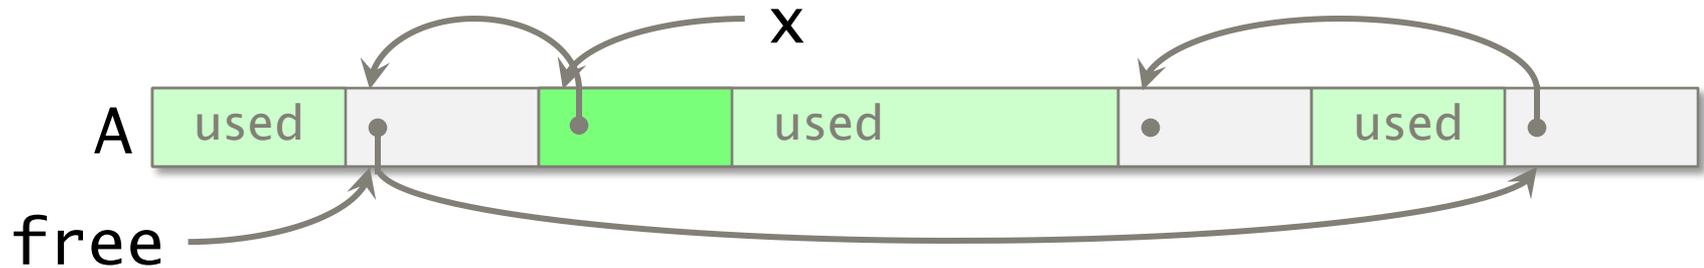
```
x = free;  
free = free->next;  
return x;
```

free object x

```
x->next = free;  
free = x;
```

Fixed-Size Deallocation

Free list



Allocate 1 object

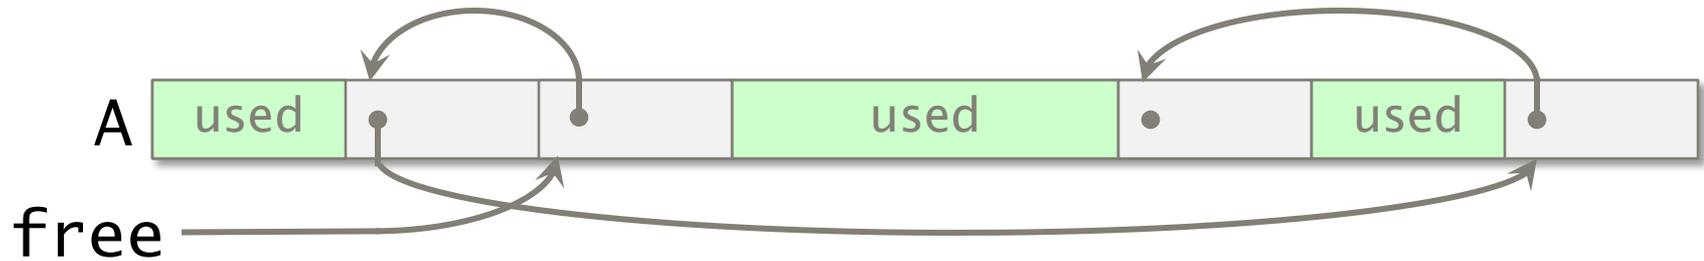
```
x = free;  
free = free->next;  
return x;
```

free object x

```
x->next = free;  
free = x;
```


Fixed-Size Deallocation

Free list



Allocate 1 object

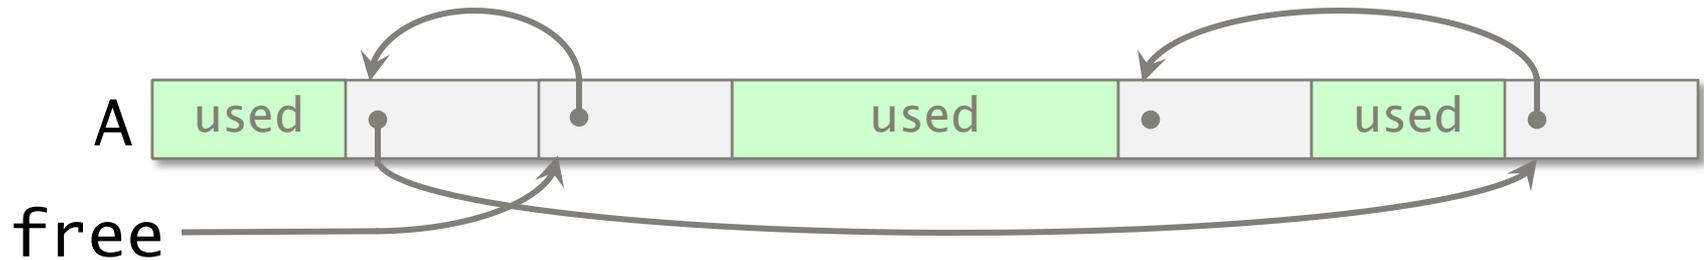
```
x = free;  
free = free->next;  
return x;
```

free object x

```
x->next = free;  
free = x;
```

Free Lists

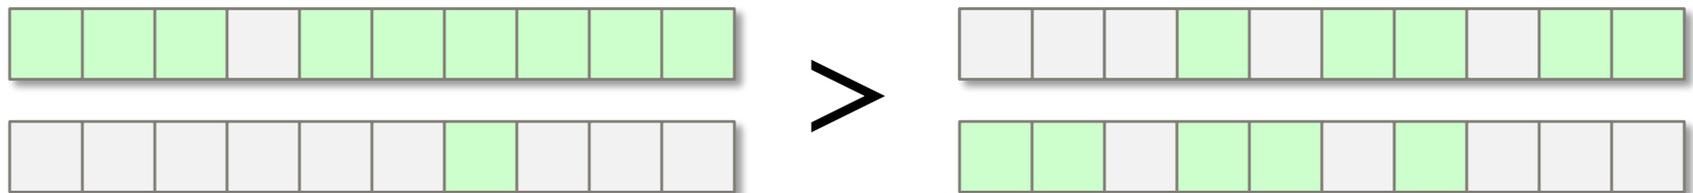
Free list



- Allocating and freeing take $\Theta(1)$ time.
- Good temporal locality.
- Poor spatial locality due to *external fragmentation* — blocks distributed across virtual memory — which can increase the size of the page table and cause disk thrashing.
- The *translation lookaside buffer (TLB)* can also be a problem.

Mitigating External Fragmentation

- Keep a free list per disk page.
- Allocate from the free list for the fullest page.
- Free a block of storage to the free list for the page on which the block resides.
- If a page becomes empty (only free-list items), the virtual-memory system can page it out without affecting program performance.
- 90-10 is better than 50-50:

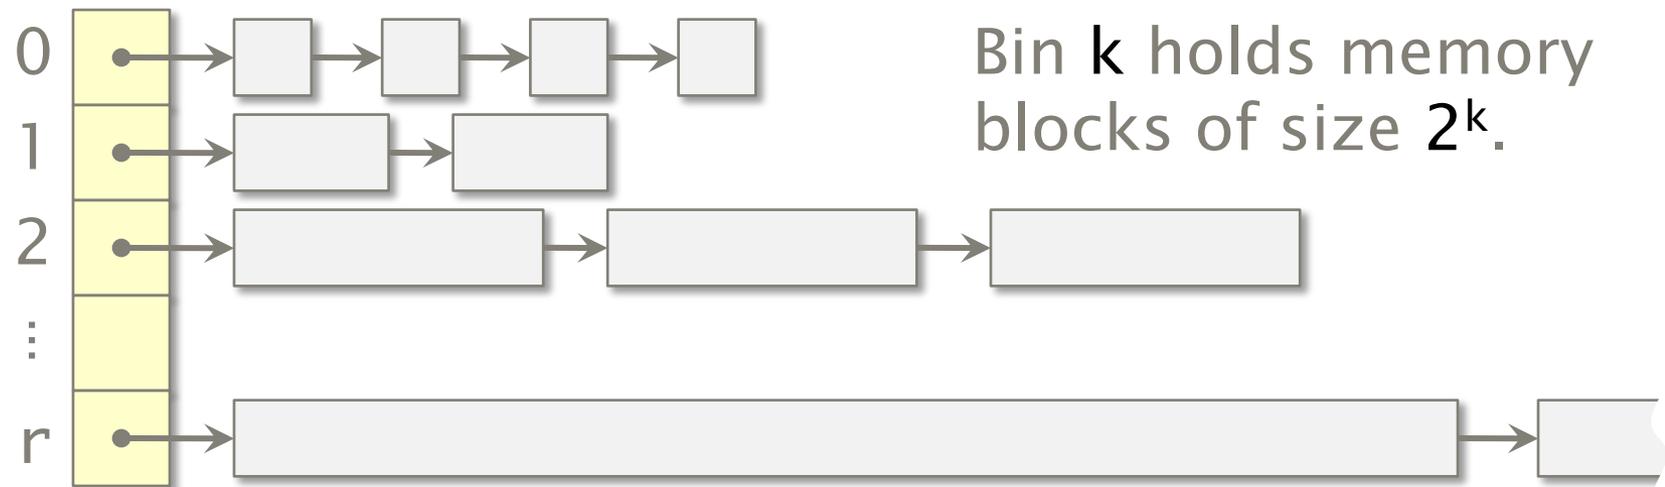


Probability that 2 random accesses hit the same page
= $.9 \times .9 + .1 \times .1 = .82$ versus $.5 \times .5 + .5 \times .5 = .5$

Variable-Sized Allocation

Binned free lists

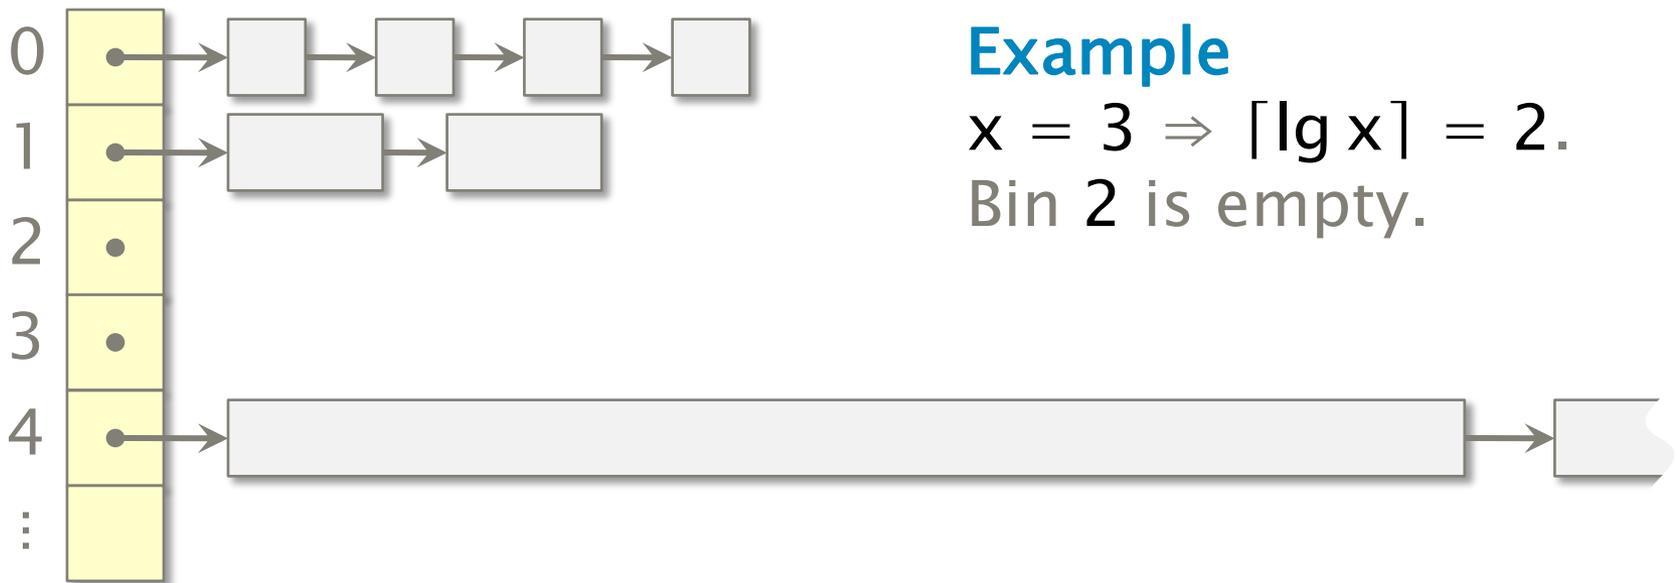
- Leverage the efficiency of free lists.
- Accept some internal fragmentation.



Allocation for Binned Free Lists

Allocate x bytes

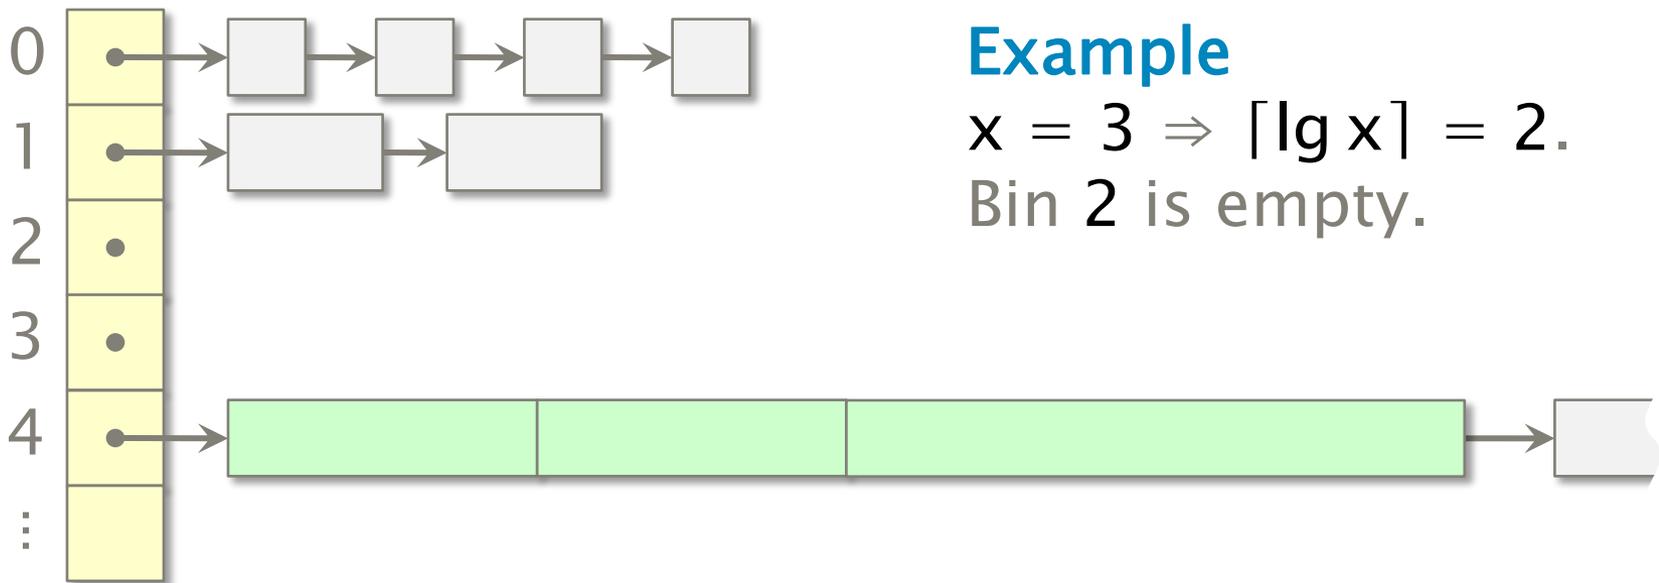
- If bin $k = \lceil \lg x \rceil$ is nonempty, return a block.
- Otherwise, find a block in the next larger nonempty bin $k' > k$, split it up into blocks of sizes $2^{k'-1}$, $2^{k'-2}$, ... 2^k , 2^k , and distribute the pieces.



Allocation for Binned Free Lists

**Allocate
x bytes**

- If bin $k = \lceil \lg x \rceil$ is nonempty, return a block.
- Otherwise, find a block in the next larger nonempty bin $k' > k$, split it up into blocks of sizes $2^{k'-1}, 2^{k'-2}, \dots, 2^k, 2^k$, and distribute the pieces.



Example

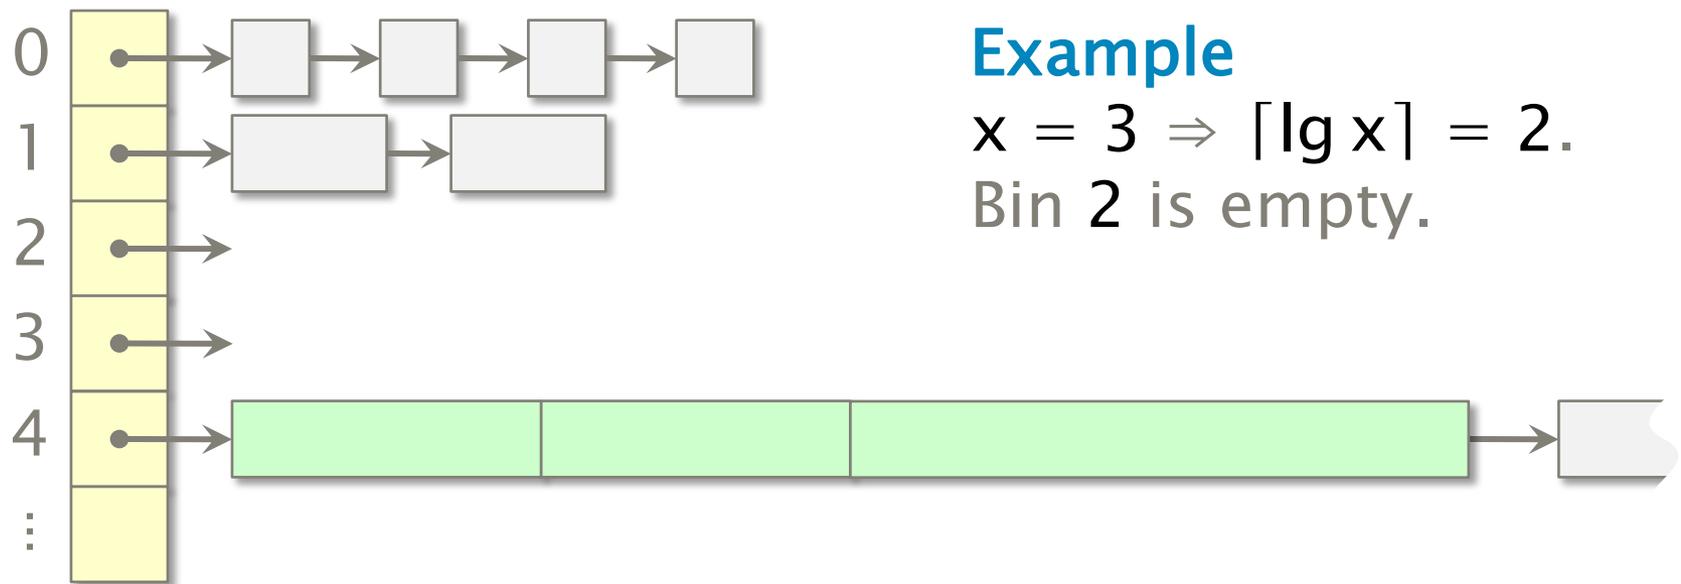
$x = 3 \Rightarrow \lceil \lg x \rceil = 2.$

Bin 2 is empty.

Allocation for Binned Free Lists

**Allocate
x bytes**

- If bin $k = \lceil \lg x \rceil$ is nonempty, return a block.
- Otherwise, find a block in the next larger nonempty bin $k' > k$, split it up into blocks of sizes $2^{k'-1}, 2^{k'-2}, \dots, 2^k, 2^k$, and distribute the pieces.



Example

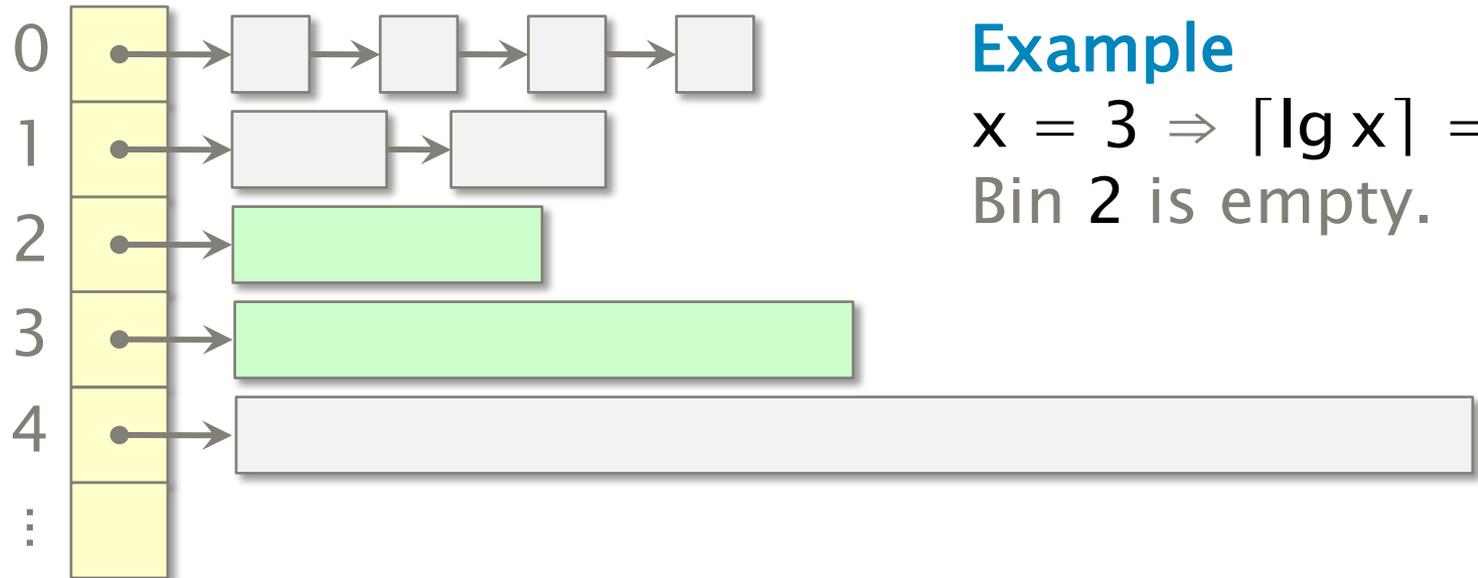
$x = 3 \Rightarrow \lceil \lg x \rceil = 2.$

Bin 2 is empty.

Allocation for Binned Free Lists

**Allocate
x bytes**

- If bin $k = \lceil \lg x \rceil$ is nonempty, return a block.
- Otherwise, find a block in the next larger nonempty bin $k' > k$, split it up into blocks of sizes $2^{k'-1}, 2^{k'-2}, \dots, 2^k, 2^k$, and distribute the pieces.*

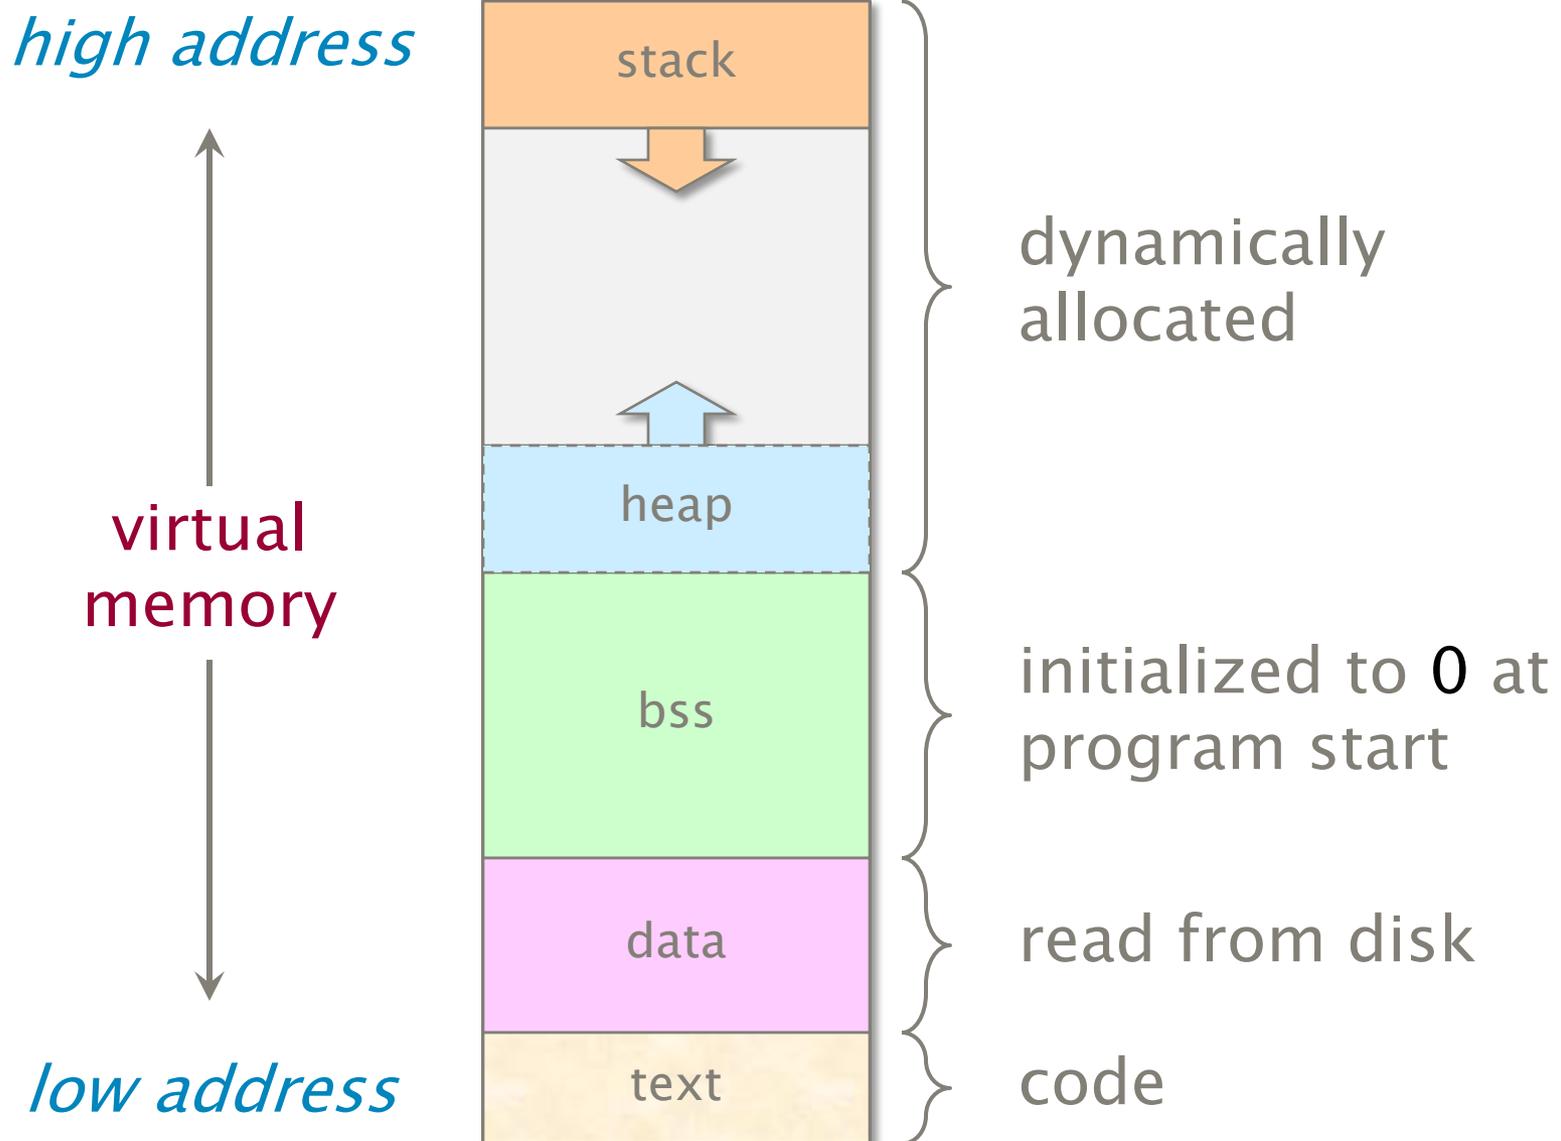


Example

$x = 3 \Rightarrow \lceil \lg x \rceil = 2$.
Bin 2 is empty.

*If no larger blocks exist, ask the OS to allocate x more bytes of VM.

Storage Layout of a Program



How Virtual is Virtual Memory?

- Q.** Since a 64-bit address space takes over a century to write at a rate of 4 billion bytes per second, we effectively never run out of virtual memory. Why not just allocate out of virtual memory and never free?
- A.** *External fragmentation* would be horrendous! The performance of the page table would degrade tremendously leading to *disk thrashing*, since all nonzero memory must be backed up on disk in page-sized blocks.

Goal of storage allocators

Use as little virtual memory as possible, and try to keep the used portions relatively compact.

Analysis of Binned Free Lists

Theorem. Suppose that the maximum amount of heap memory in use at any time by a program is M . If the heap is managed by a BFL allocator, the amount of virtual memory consumed by heap storage is $O(M \lg M)$.

Proof. An allocation request for a block of size x consumes $2^{\lceil \lg x \rceil} \leq 2x$ storage. Thus, the amount of virtual memory devoted to blocks of size 2^k is at most $2M$. Since there are at most $\lg M$ free lists, the theorem holds. ■

⇒ In fact, BFL is $\Theta(1)$ -competitive with the optimal allocator (assuming no coalescing).

Coalescing

Binned free lists can sometimes be heuristically improved by splicing together adjacent small blocks into a larger block.

- Clever schemes exist for finding adjacent blocks efficiently, e.g., the “buddy” system, but the overhead is still greater than simple BFL.
- No good theoretical bounds exist that prove the effectiveness of coalescing.
- Coalescing seems to work in practice, because storage tends to be deallocated as a stack (LIFO) or in batches.

Garbage Collectors

Idea

- Free the programmer from freeing objects.
- A garbage collector identifies and recycles the objects that the program can no longer access.
- GC can be built-in (Java, Python) or do-it-yourself.



Garbage Collection

Terminology

Roots are objects directly accessible by the program (globals, stack, etc.).

Live objects are reachable from the roots by following pointers.

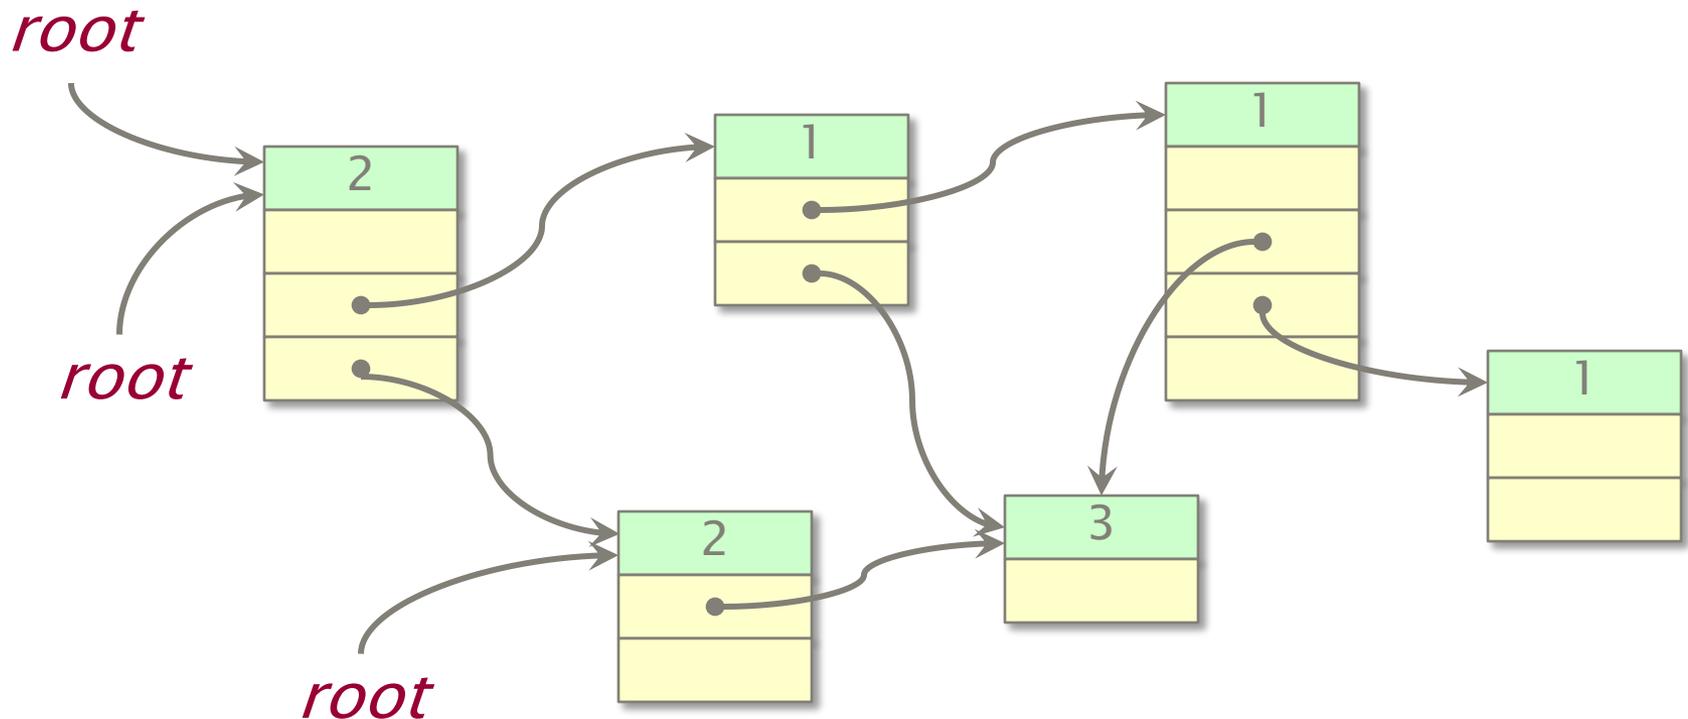
Dead objects are inaccessible and can be recycled.

How can the GC identify pointers in objects?

- Strong typing.
- Prohibit pointer arithmetic (which may slow down some programs).

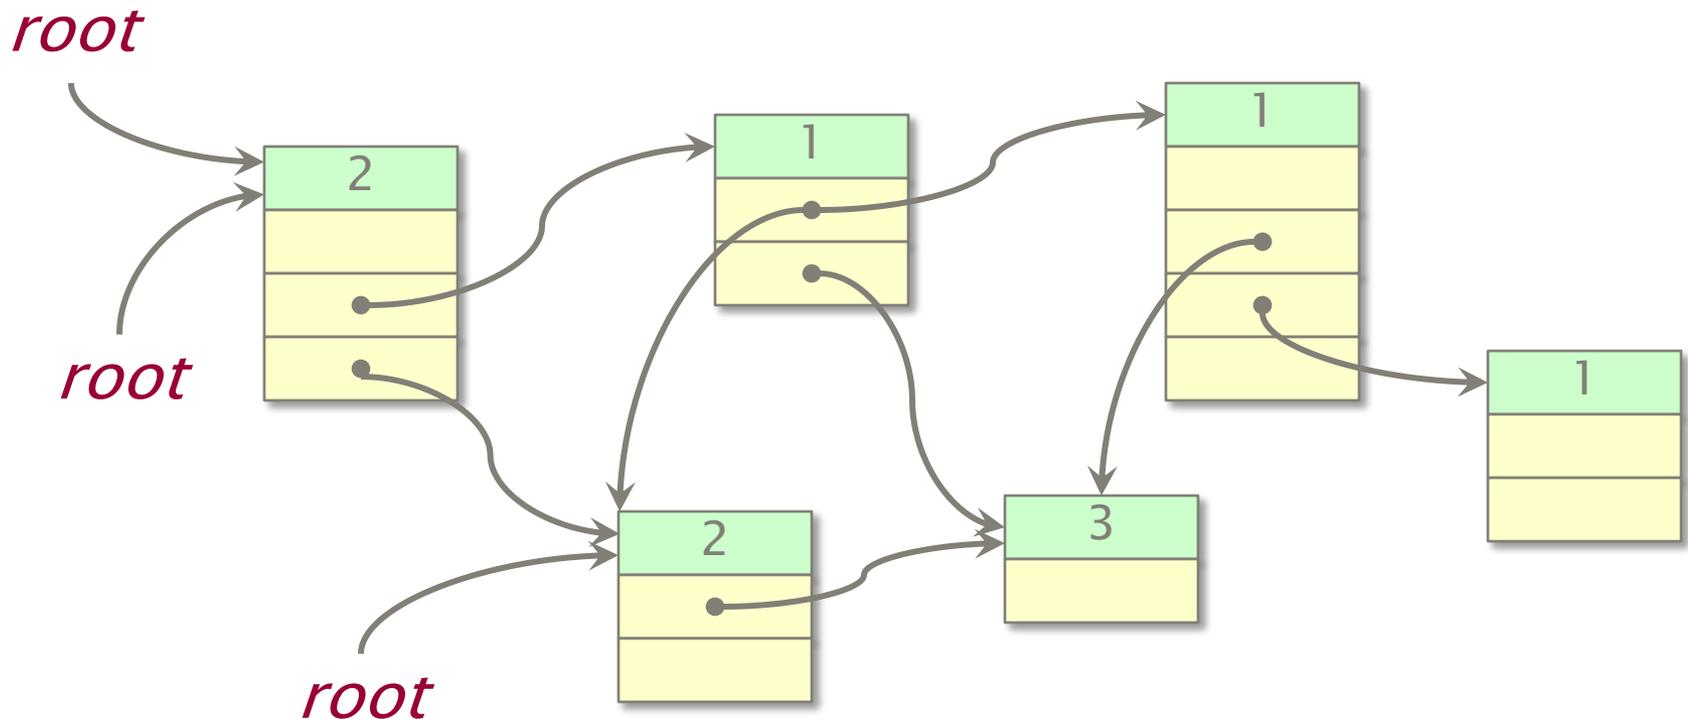
Reference Counting

Keep a count of the number of pointers referencing each object. If the count drops to 0, free the dead object.



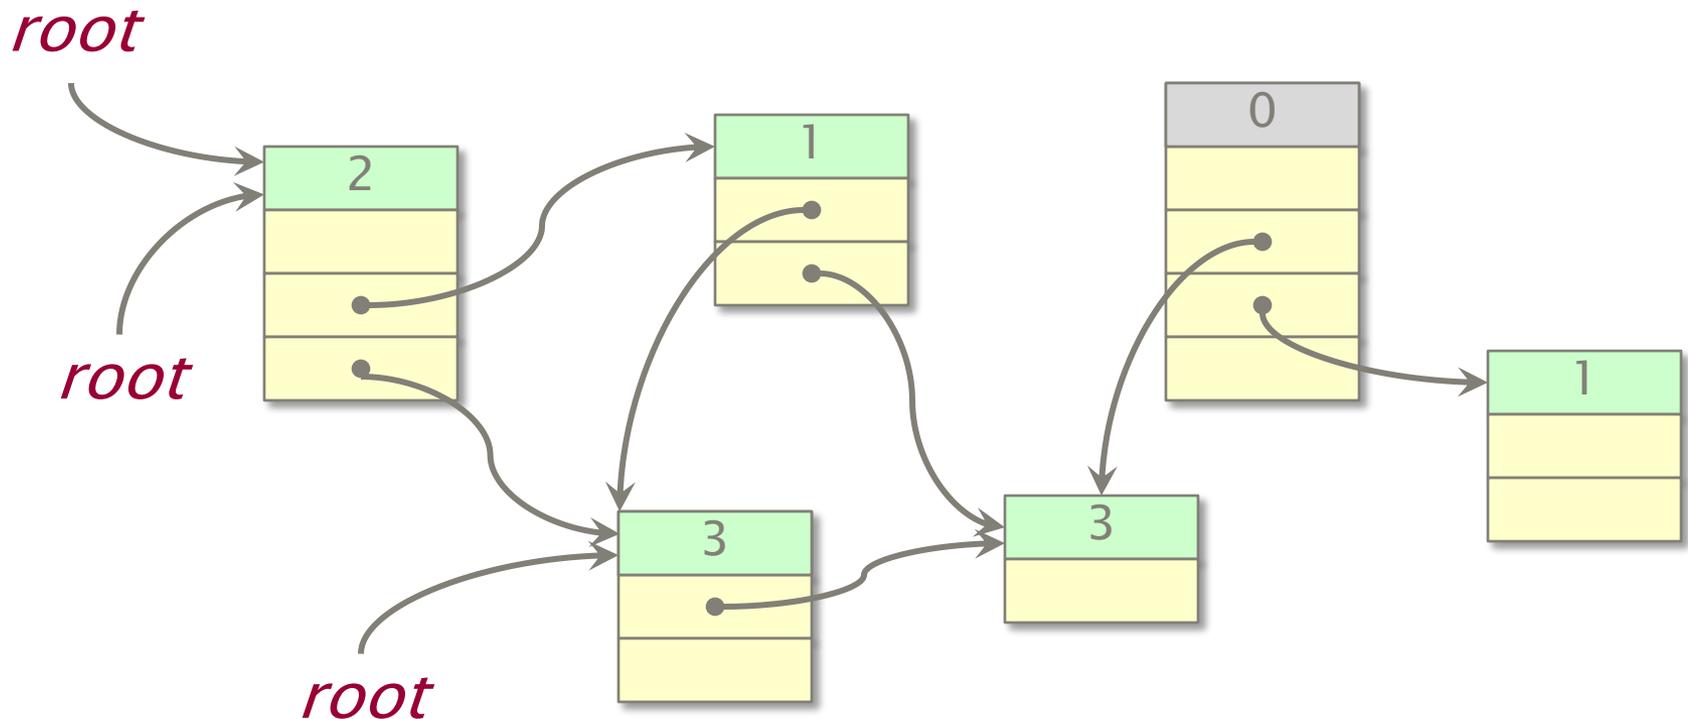
Reference Counting

Keep a count of the number of pointers referencing each object. If the count drops to 0, free the dead object.



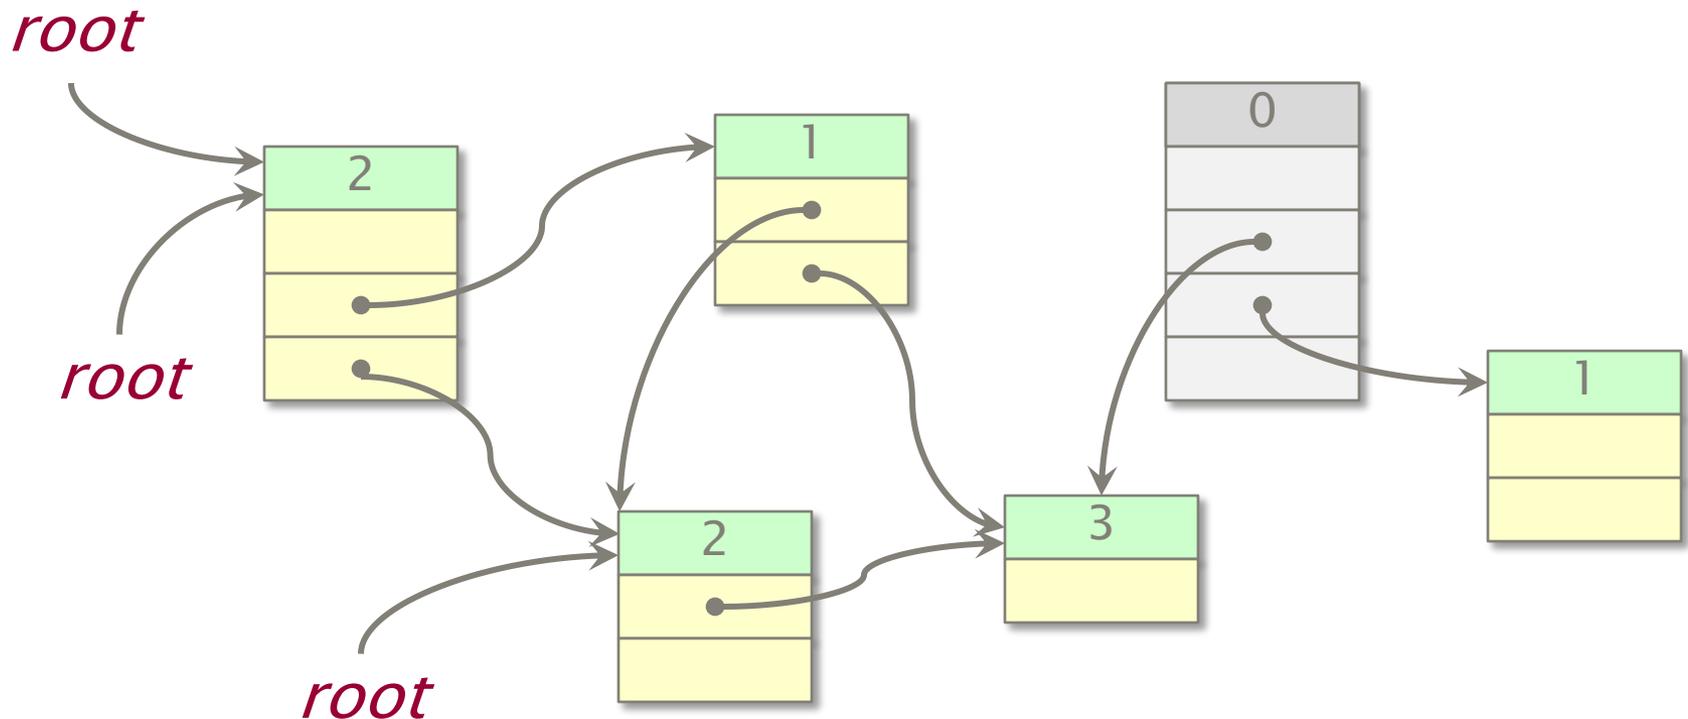
Reference Counting

Keep a count of the number of pointers referencing each object. If the count drops to 0, free the dead object.



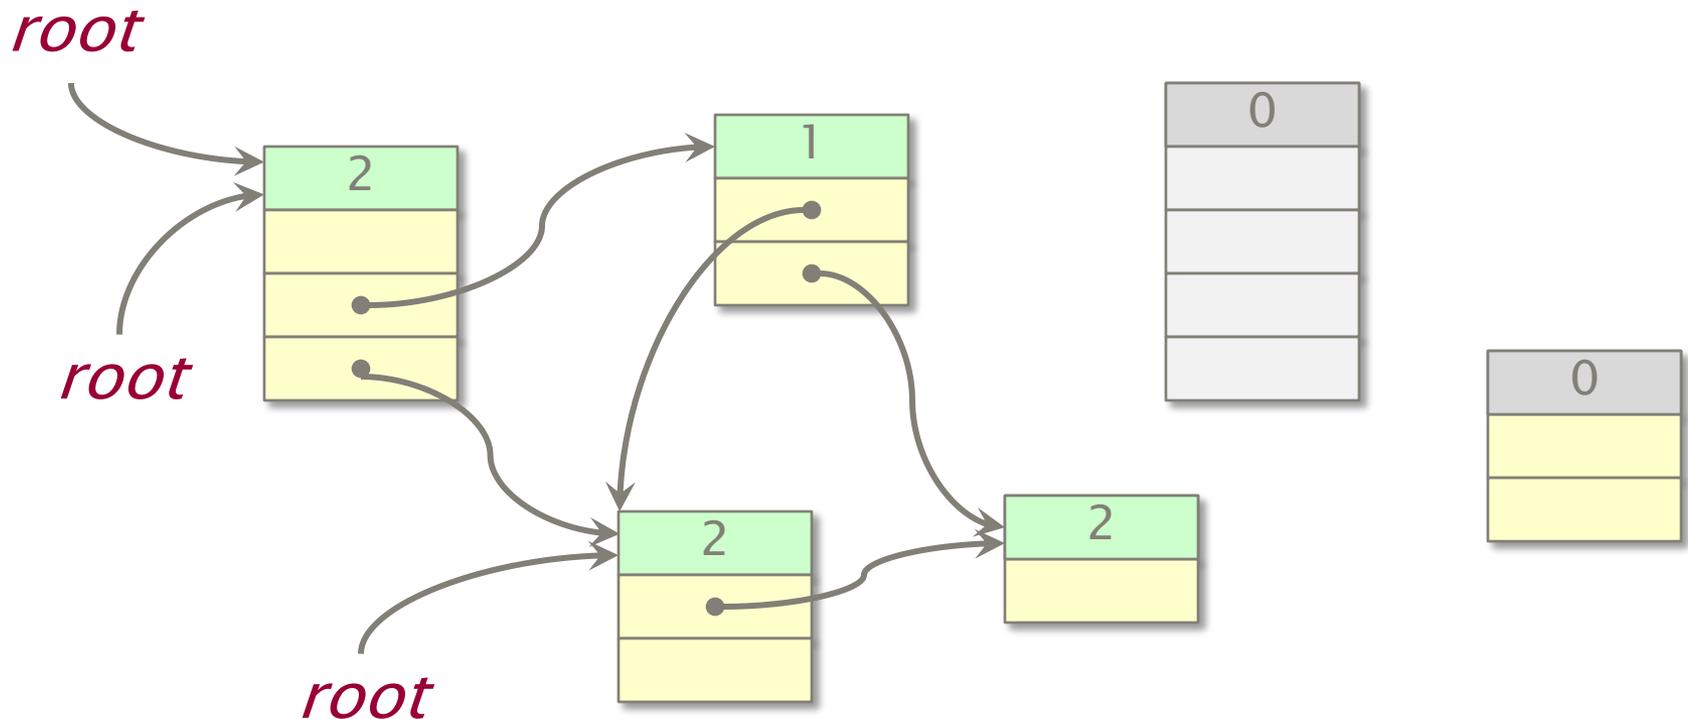
Reference Counting

Keep a count of the number of pointers referencing each object. If the count drops to 0, free the dead object.



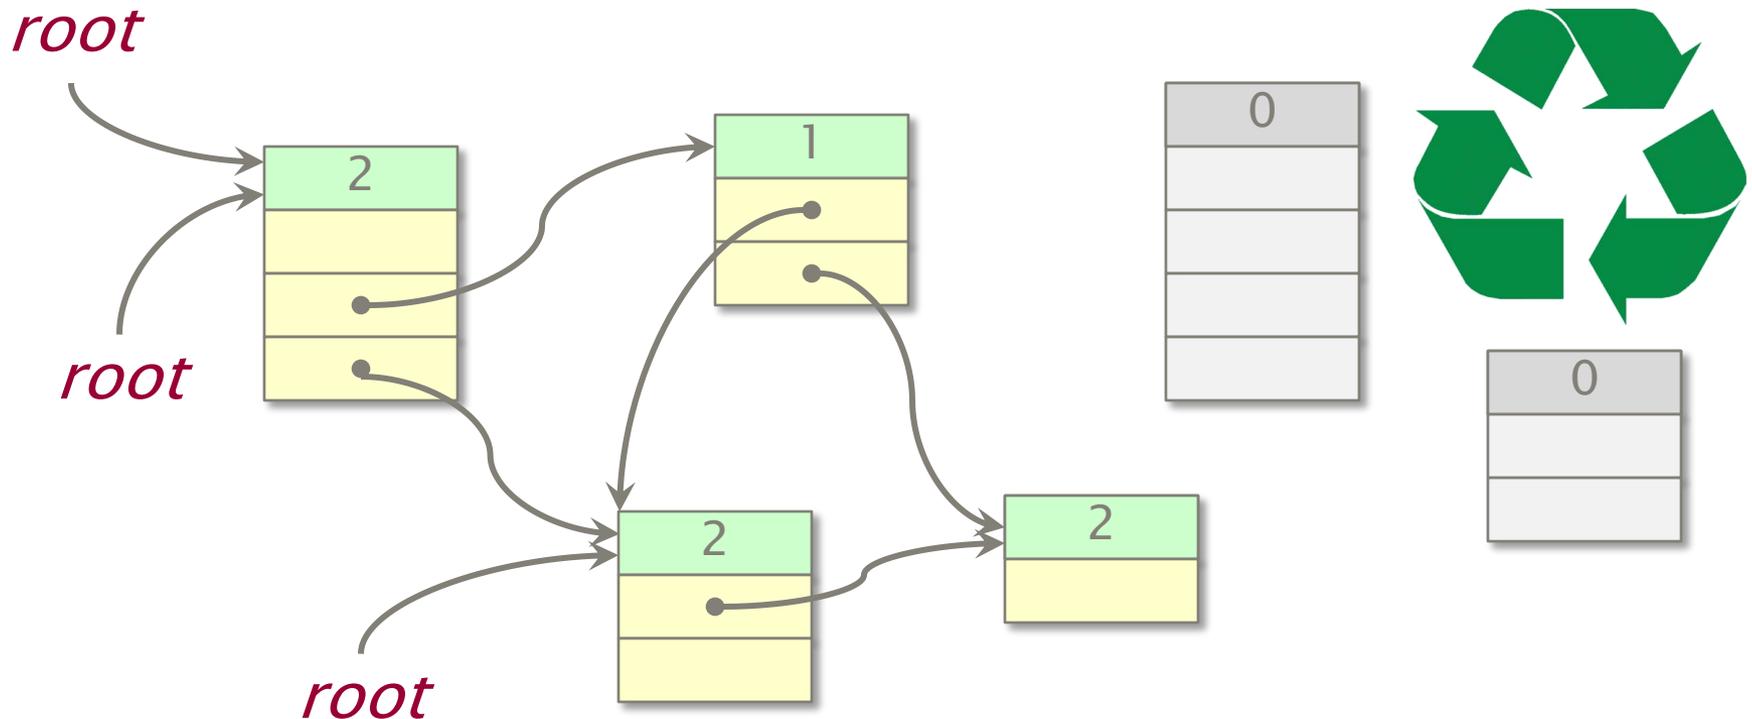
Reference Counting

Keep a count of the number of pointers referencing each object. If the count drops to 0, free the dead object.



Reference Counting

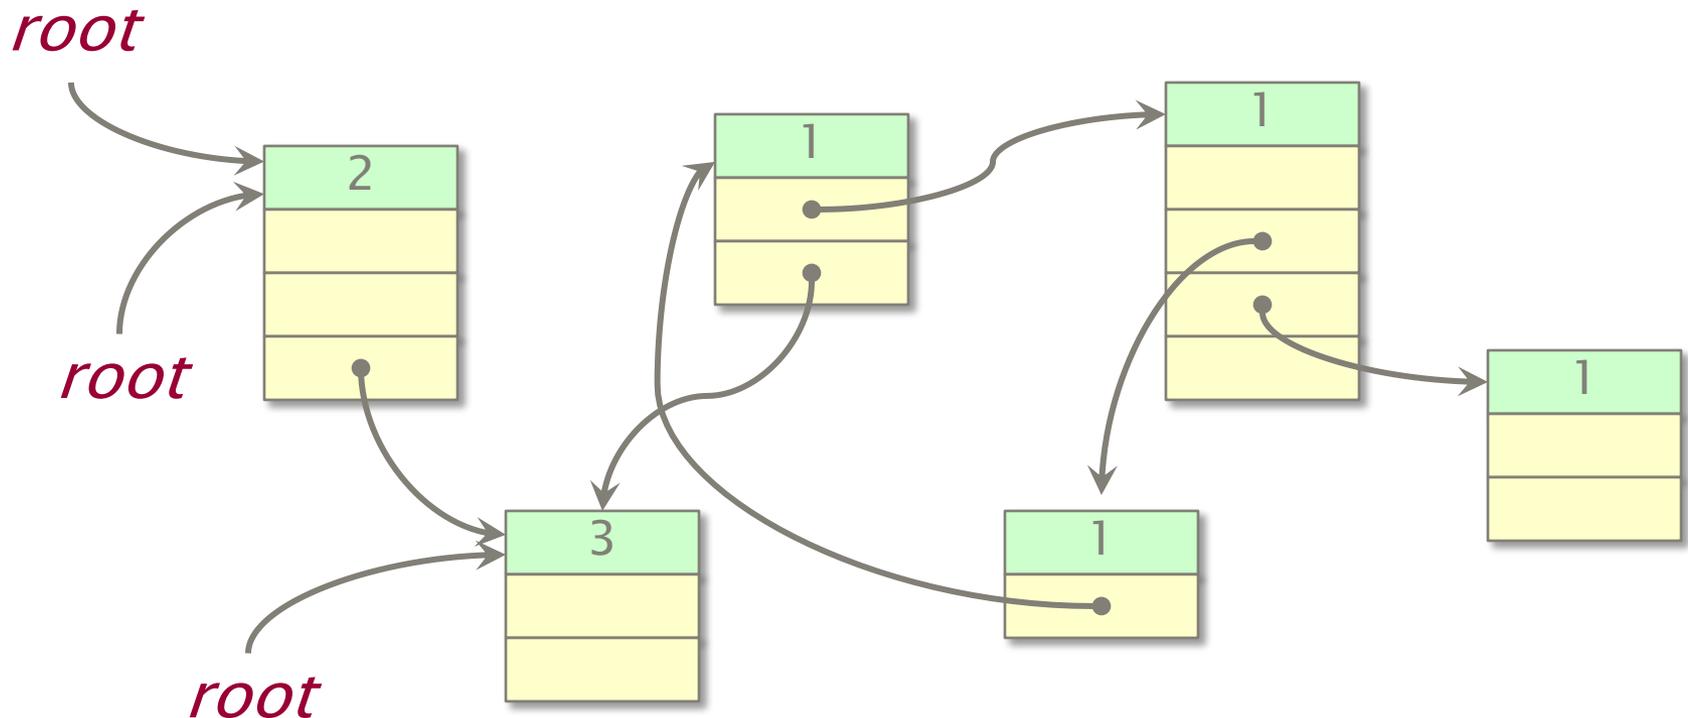
Keep a count of the number of pointers referencing each object. If the count drops to 0, free the dead object.



Limitation of Reference Counting

Problem

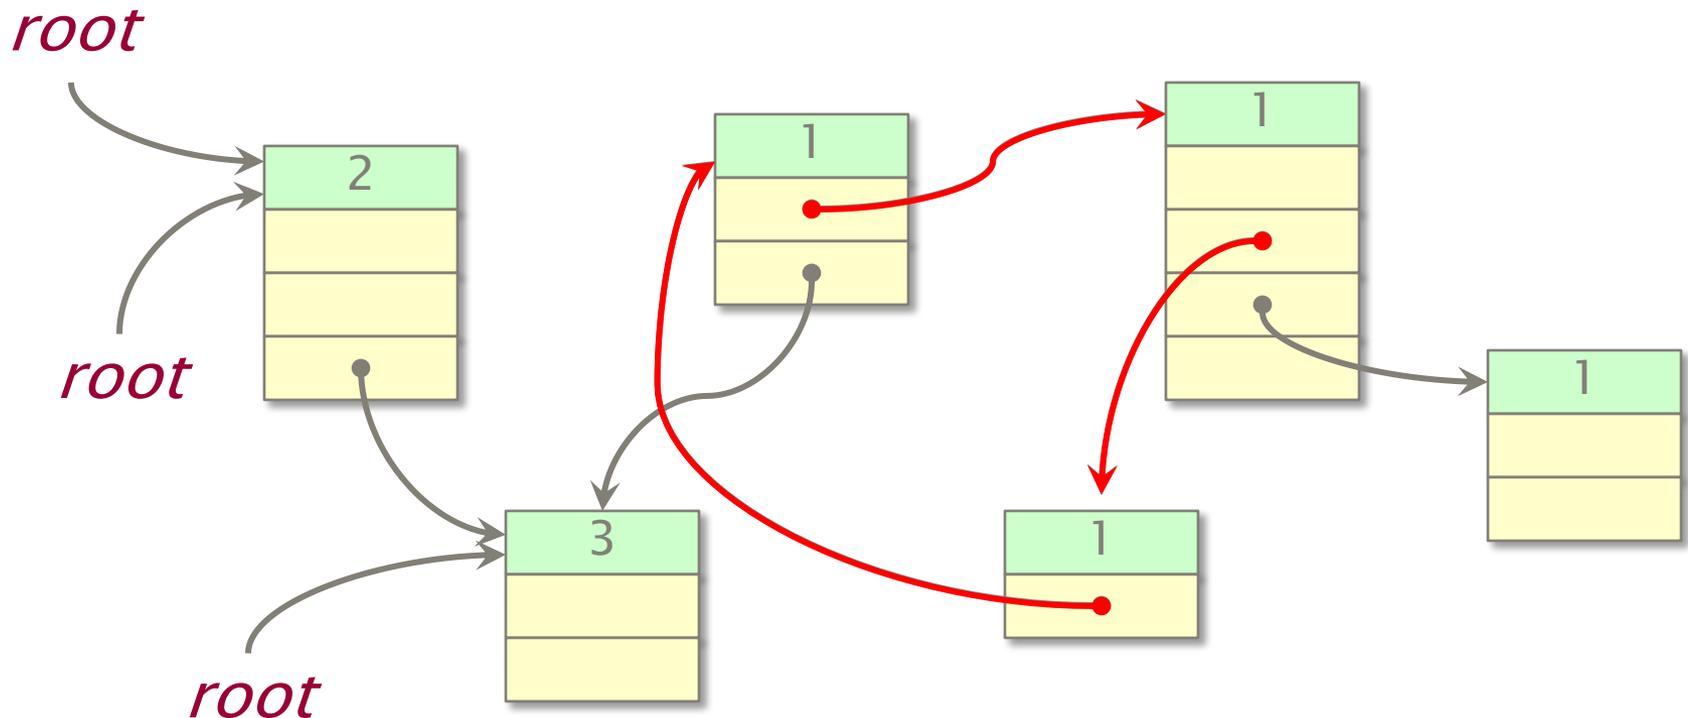
A cycle is never garbage collected!



Limitation of Reference Counting

Problem

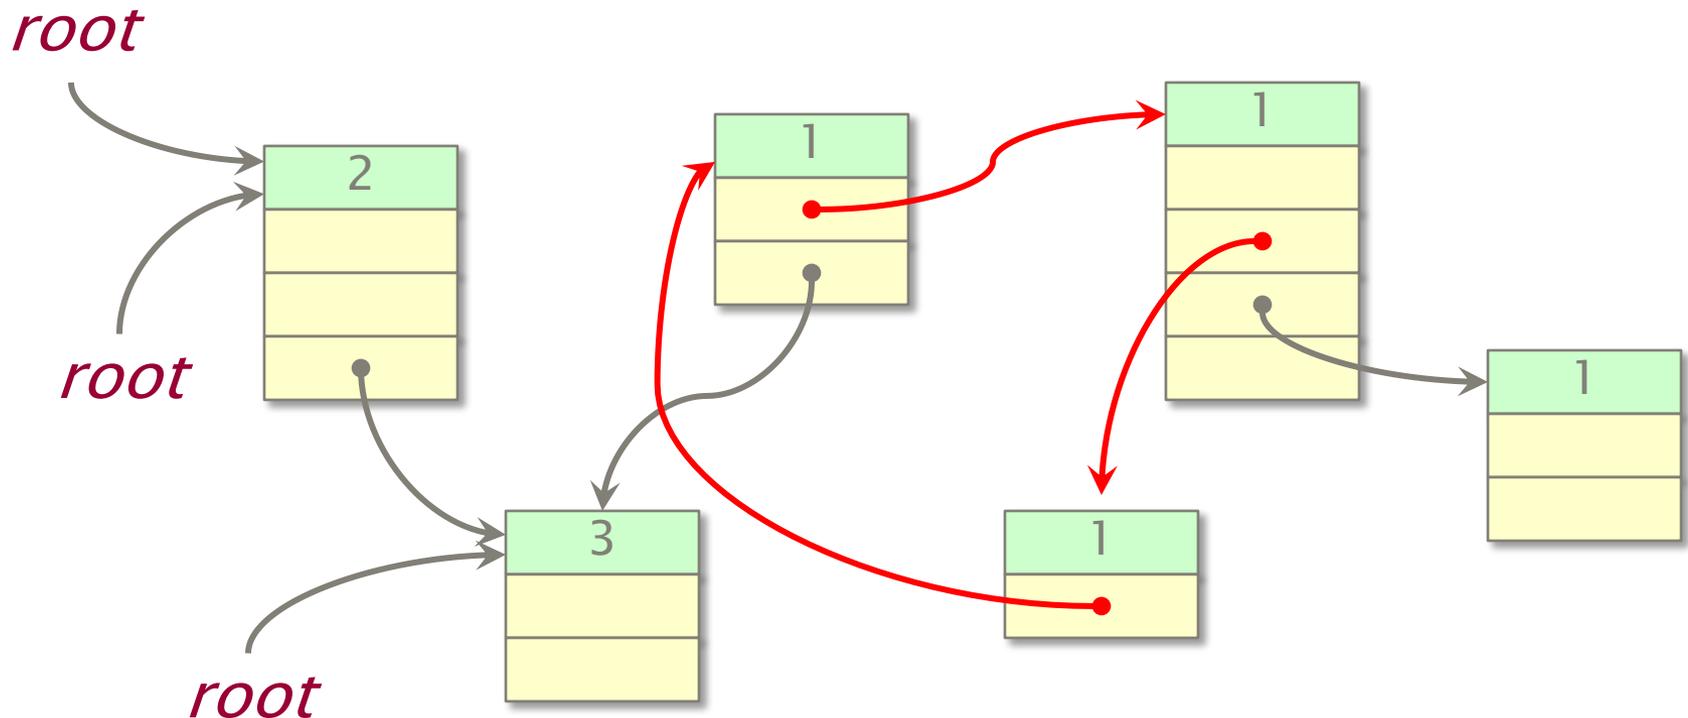
A cycle is never garbage collected!



Limitation of Reference Counting

Problem

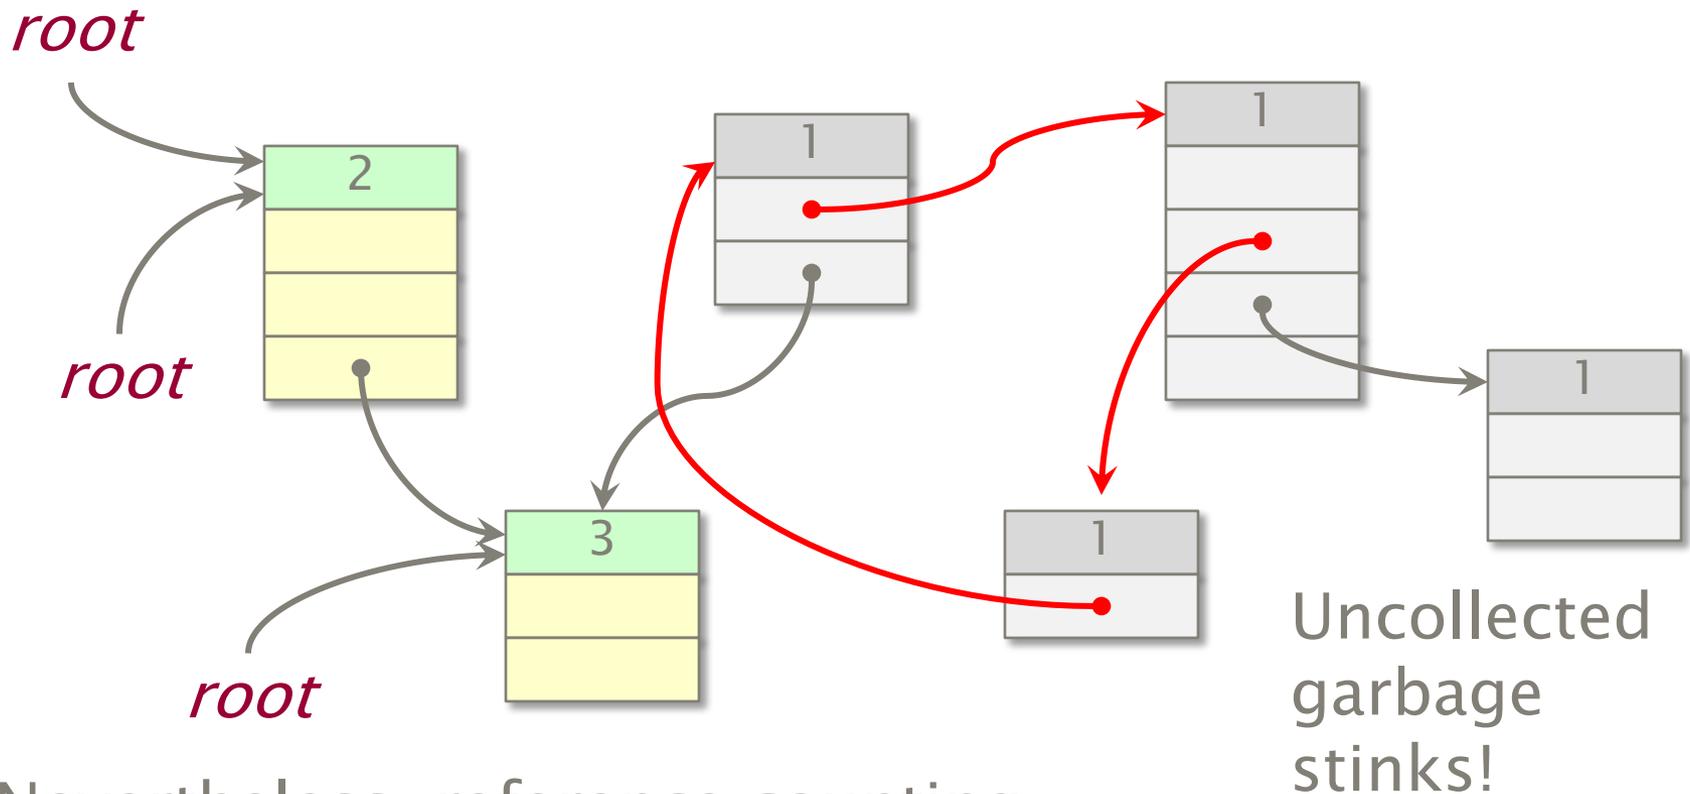
A cycle is never garbage collected!



Limitation of Reference Counting

Problem

A cycle is never garbage collected!



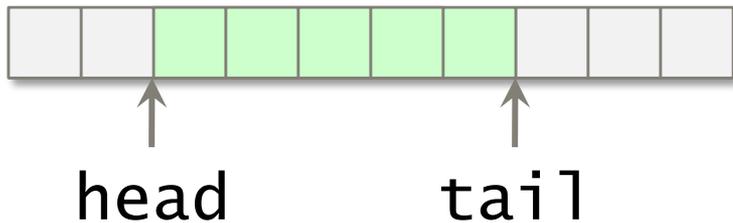
Nevertheless, reference counting works well for acyclic structures.

Graph Abstraction

Idea

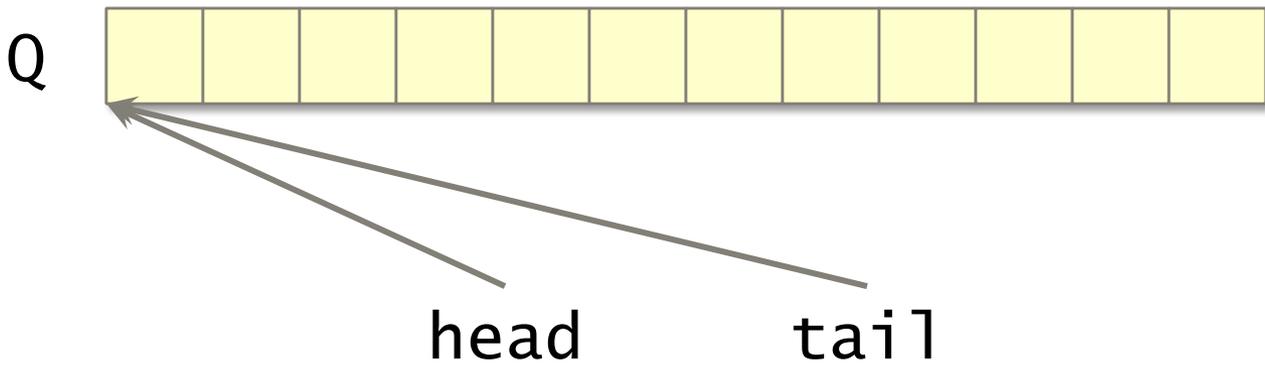
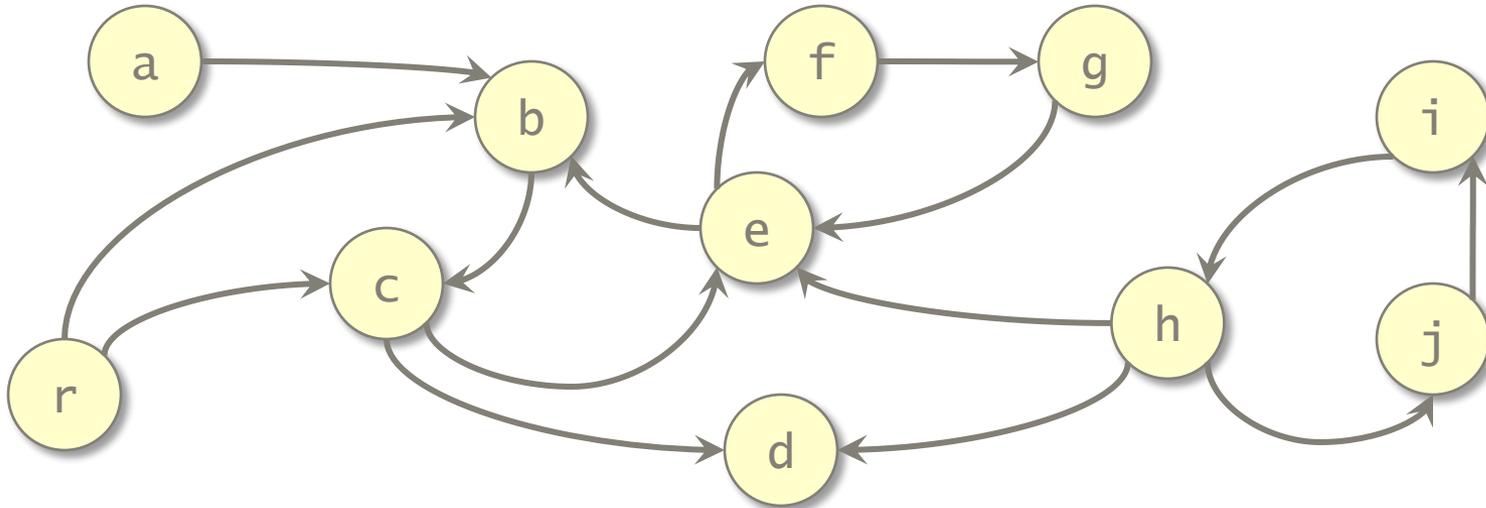
Objects and pointers form a directed graph $G = (V, E)$. Live objects are reachable from the roots. Use breadth-first search to find the live objects.

FIFO queue Q

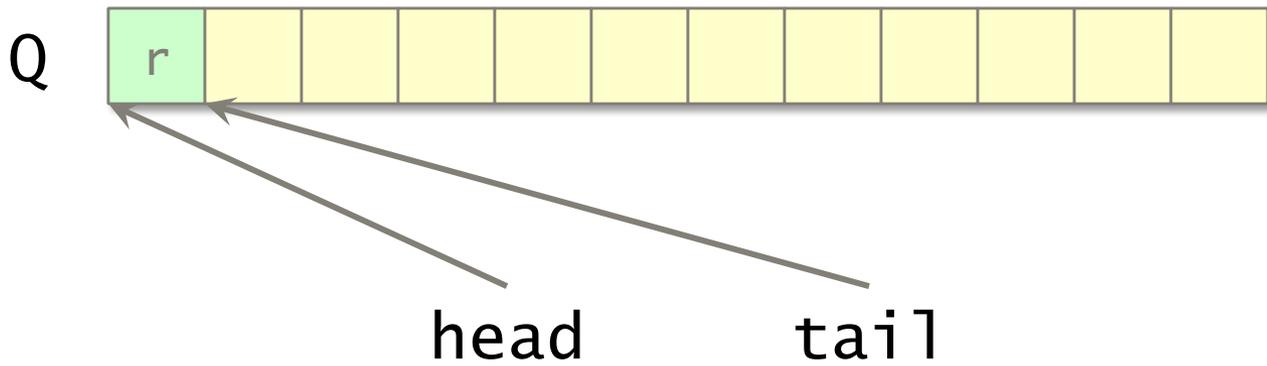
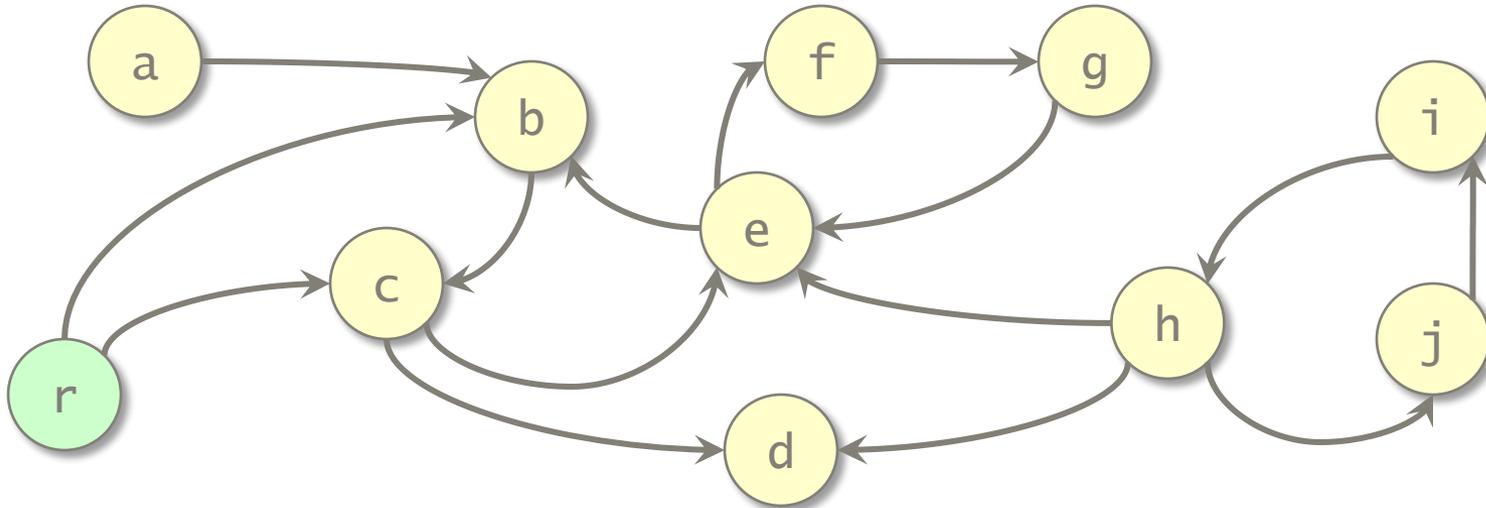


```
for ( $\forall v \in V$ ) {  
  if (root(v)) {  
    v.mark = 1;  
    enqueue(Q, v);  
  } else v.mark = 0;  
  
while (Q  $\neq \emptyset$ ) {  
  u = dequeue(Q);  
  for ( $\forall v \in V$  such that  $(u, v) \in E$ ) {  
    if (v.mark == 0) {  
      v.mark = 1;  
      enqueue(Q, v);  
    }  
  }  
}
```

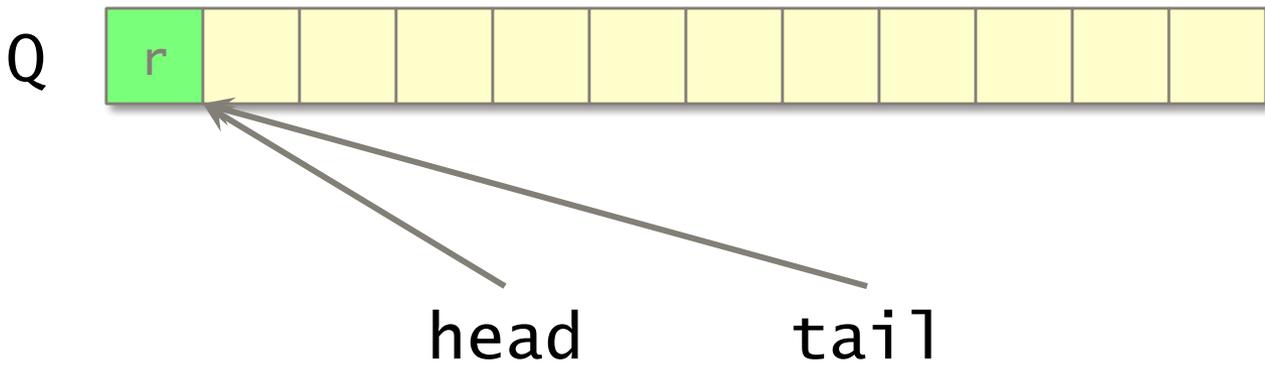
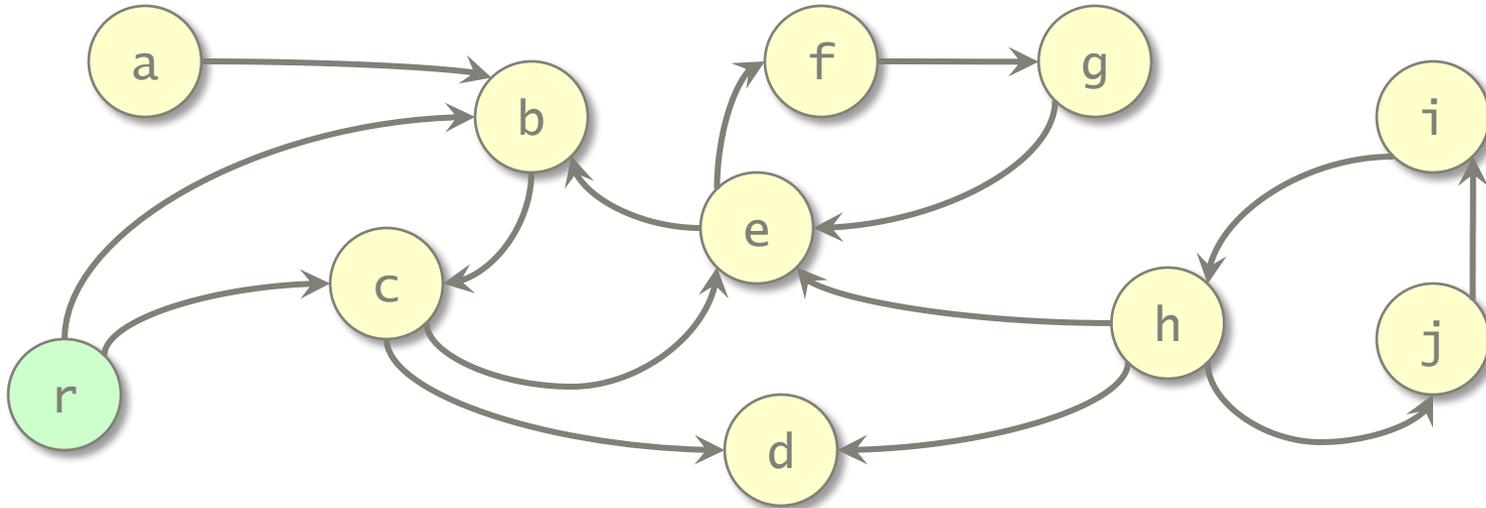
Breadth-First Search



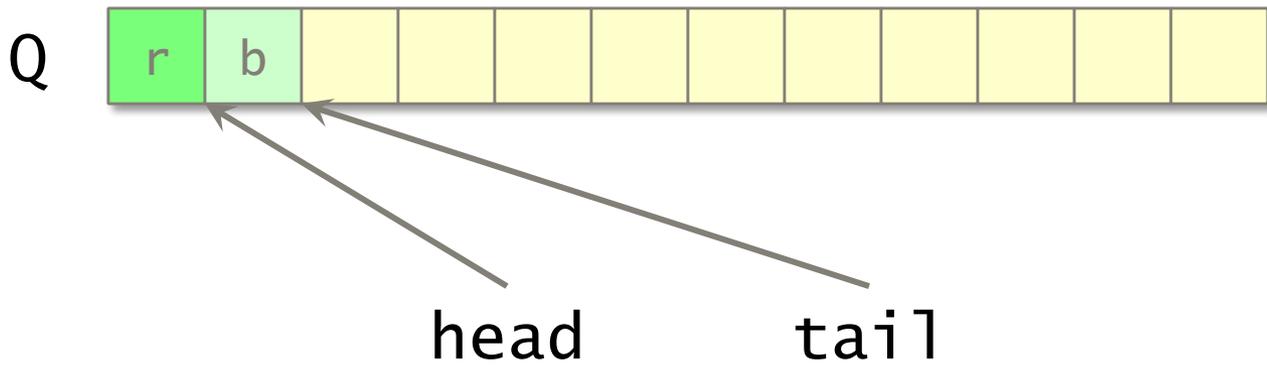
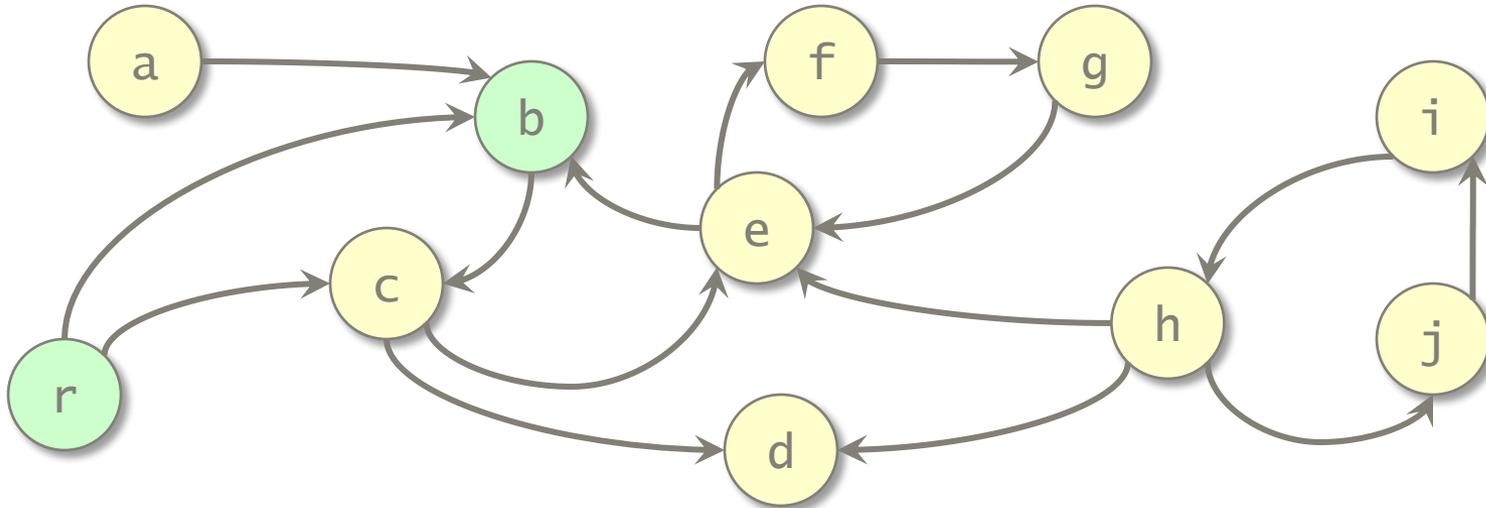
Breadth-First Search



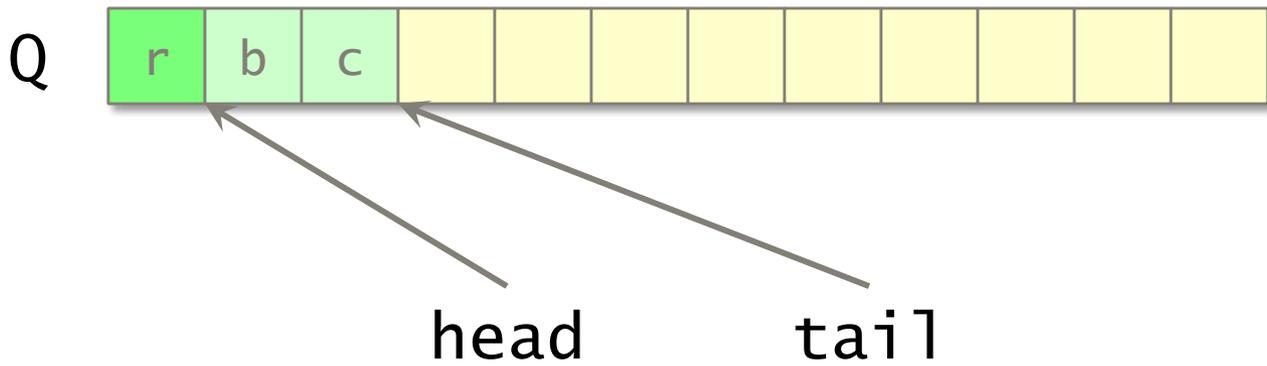
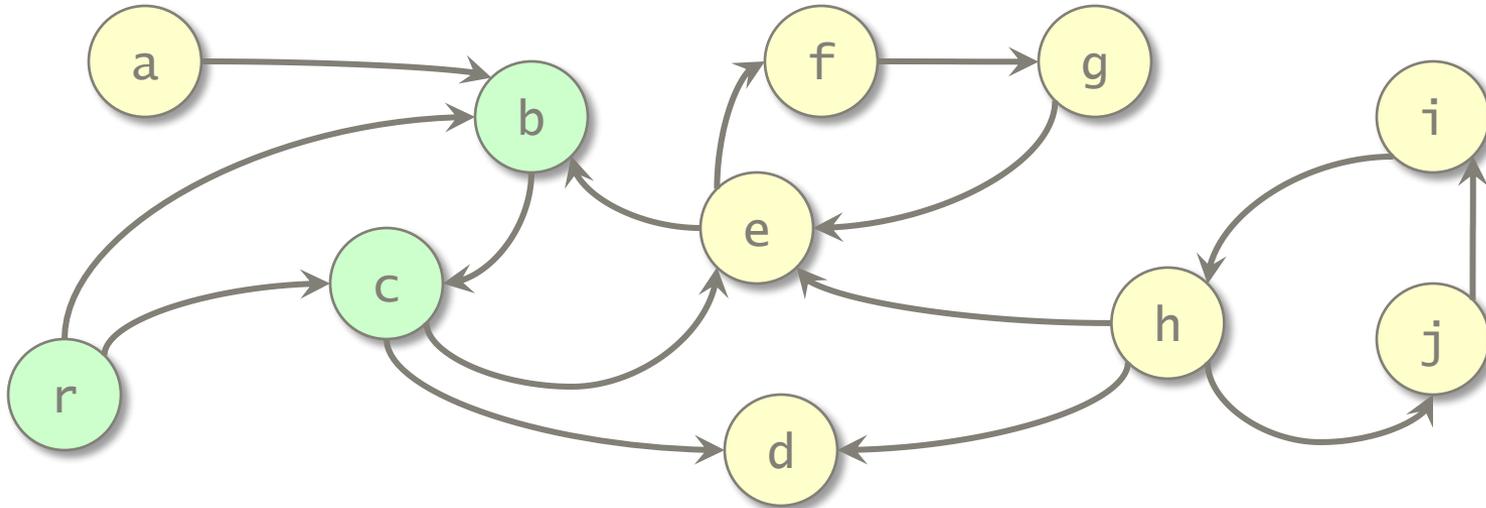
Breadth-First Search



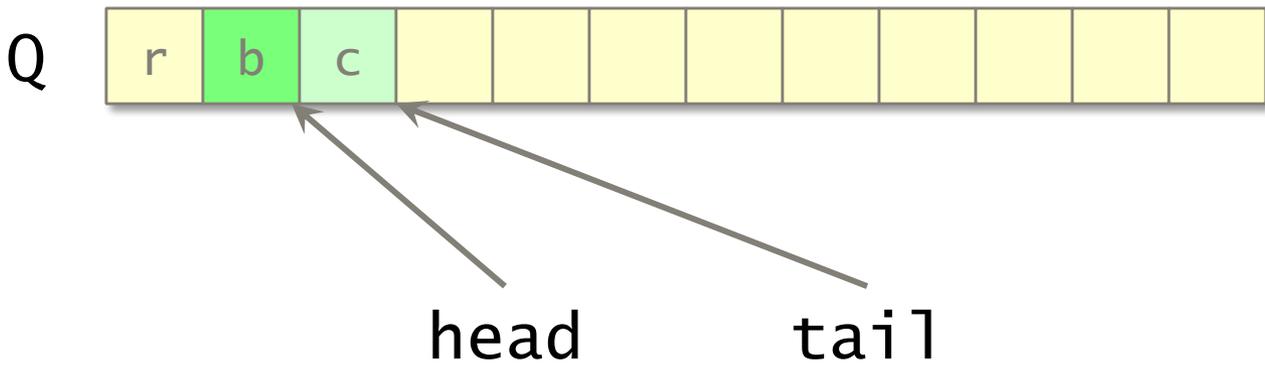
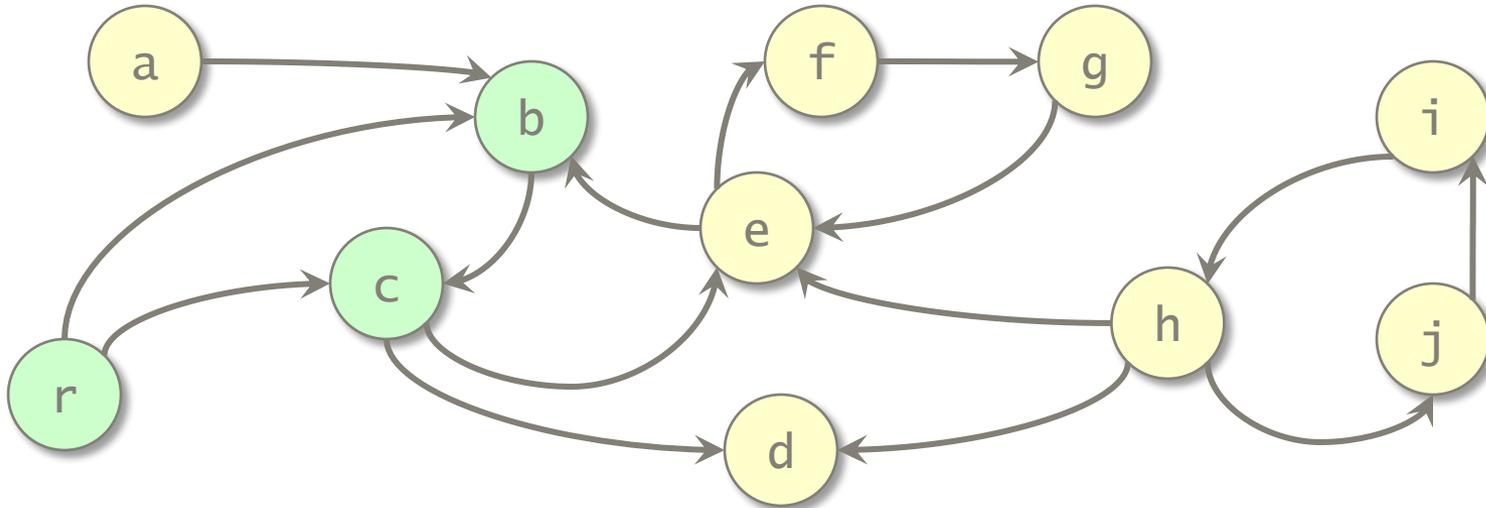
Breadth-First Search



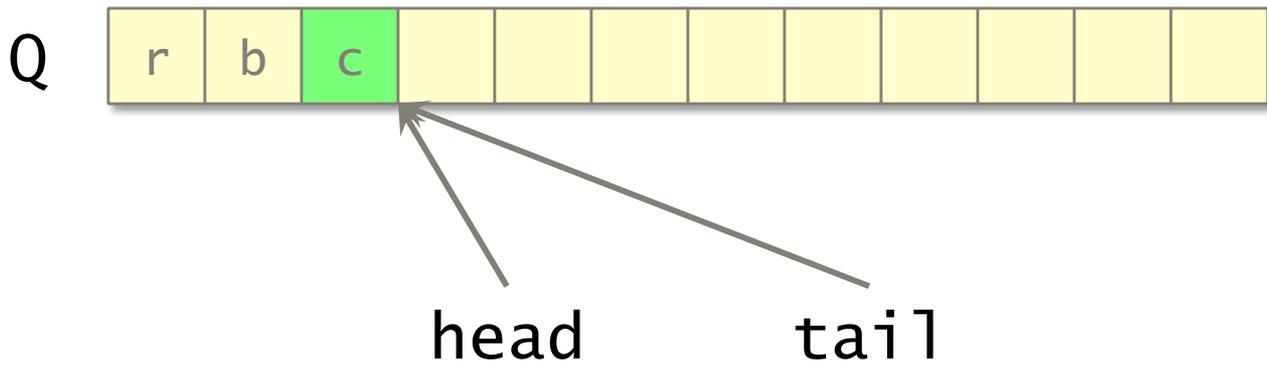
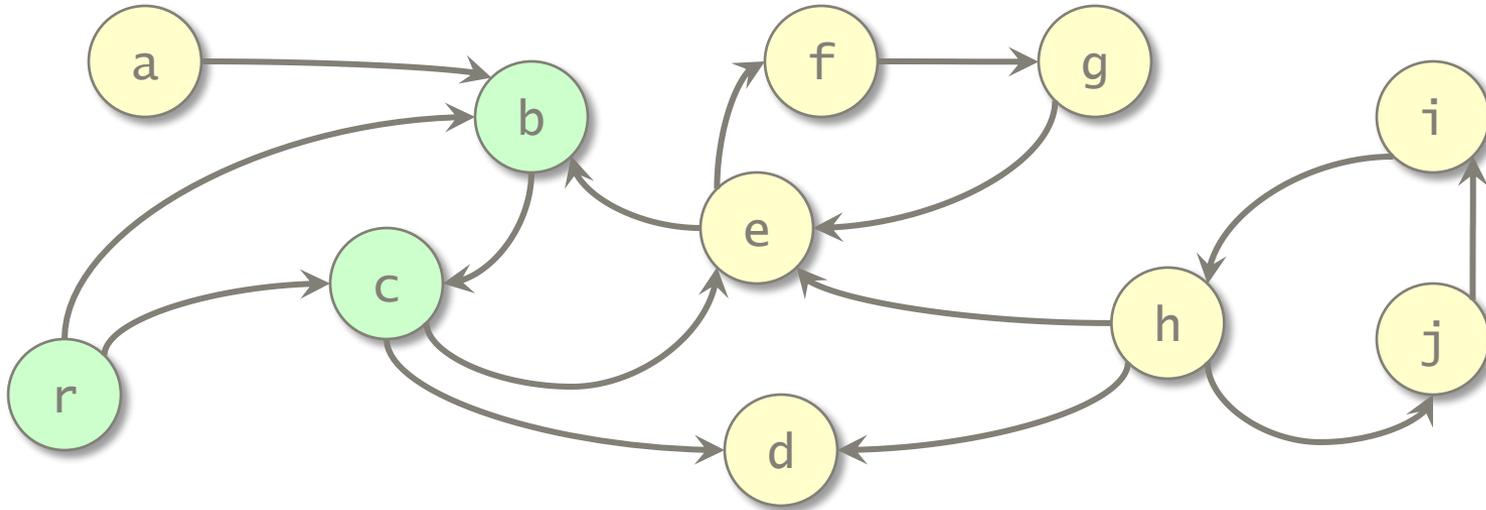
Breadth-First Search



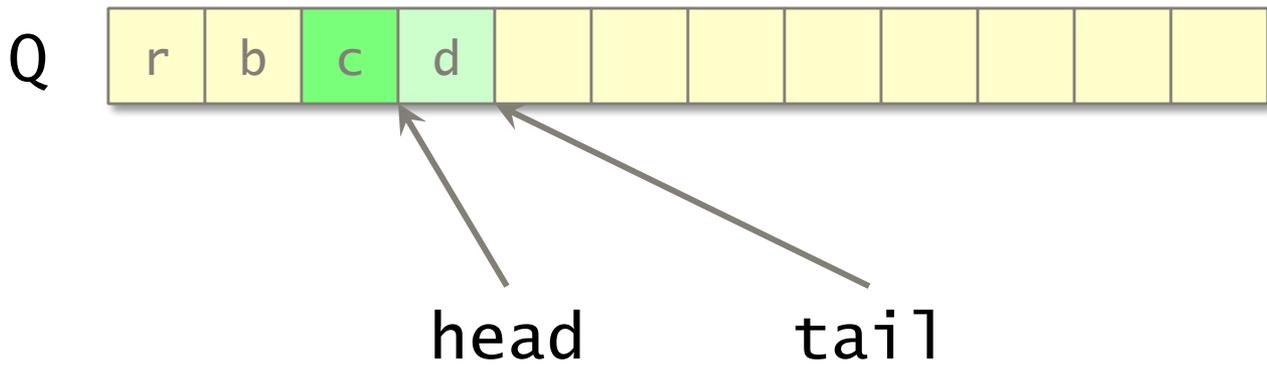
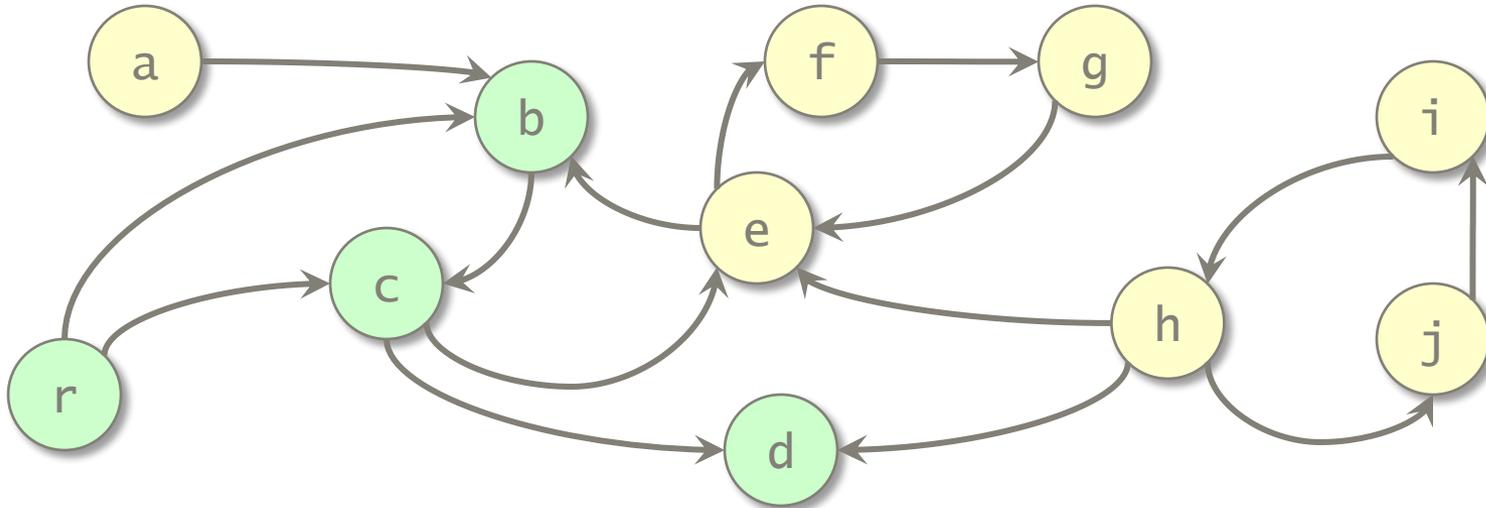
Breadth-First Search



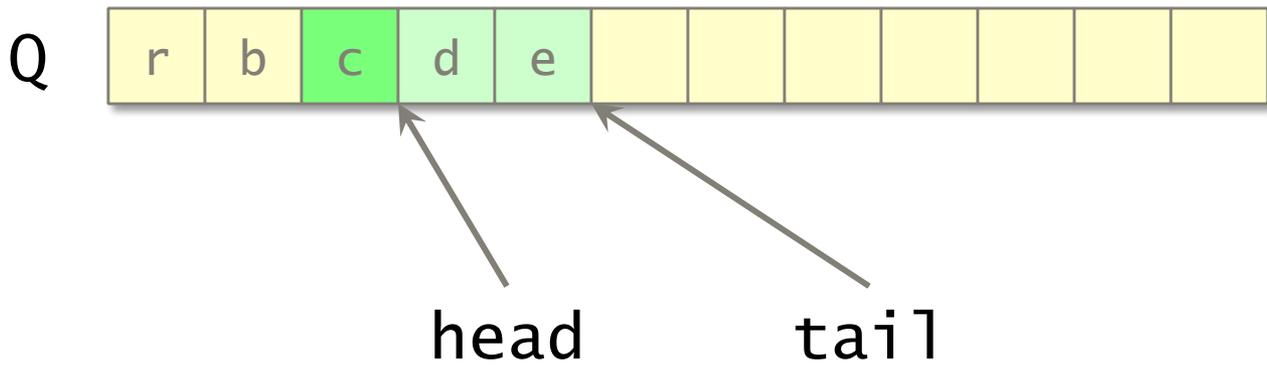
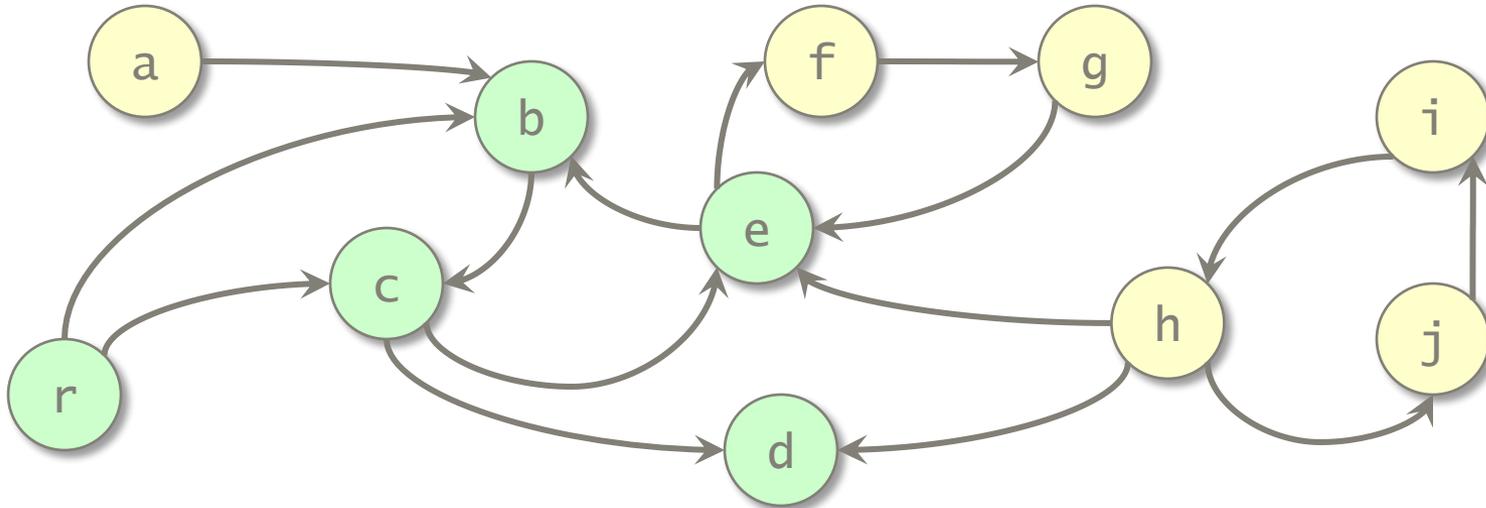
Breadth-First Search



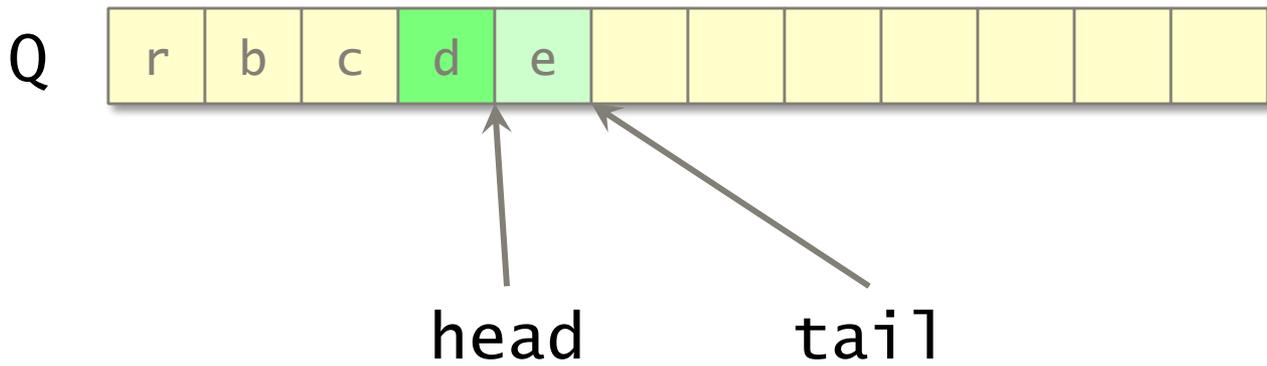
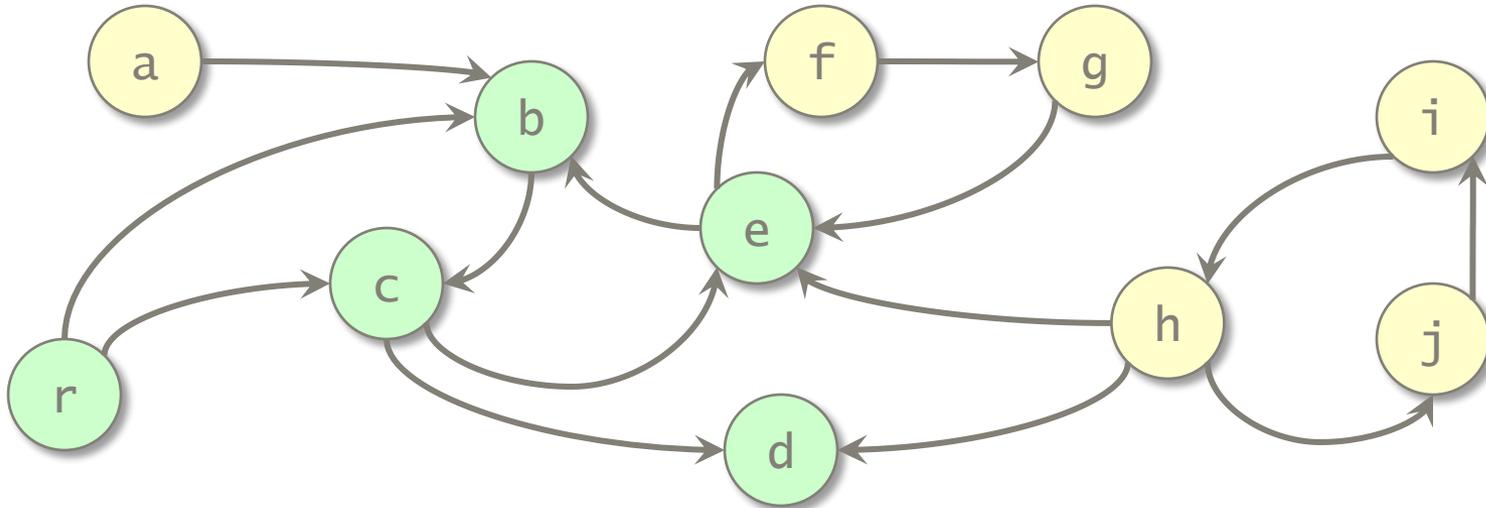
Breadth-First Search



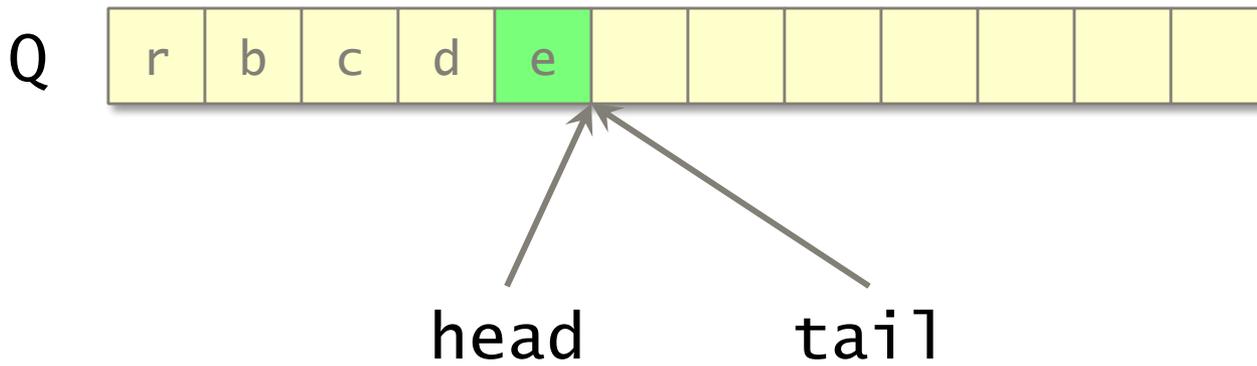
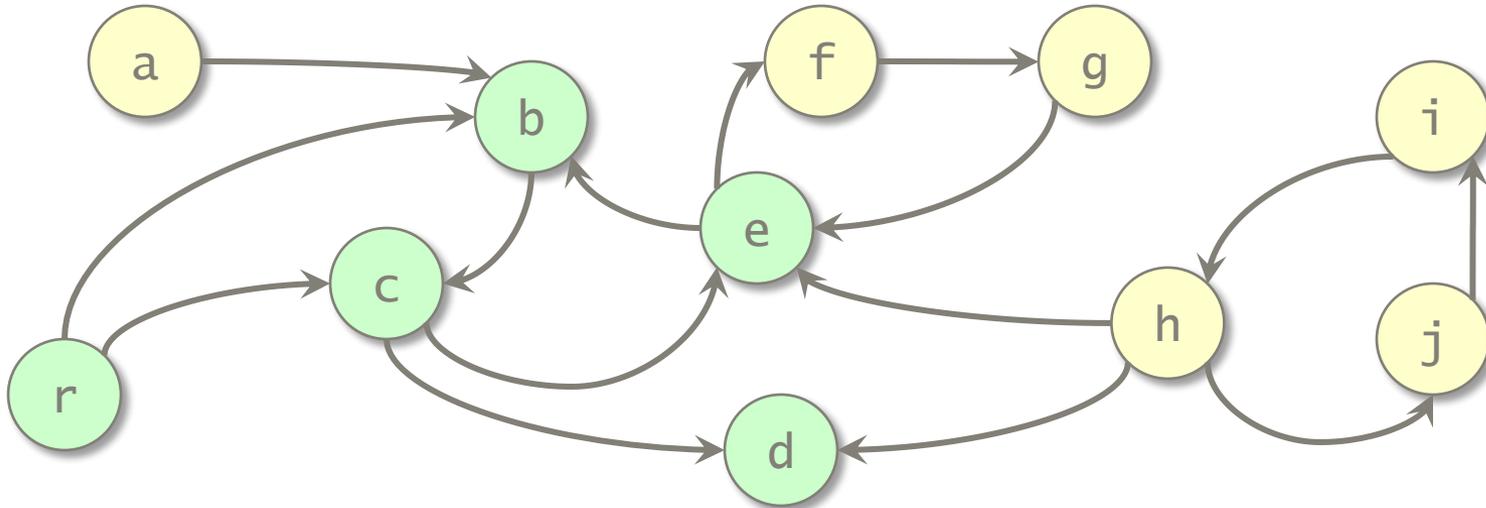
Breadth-First Search



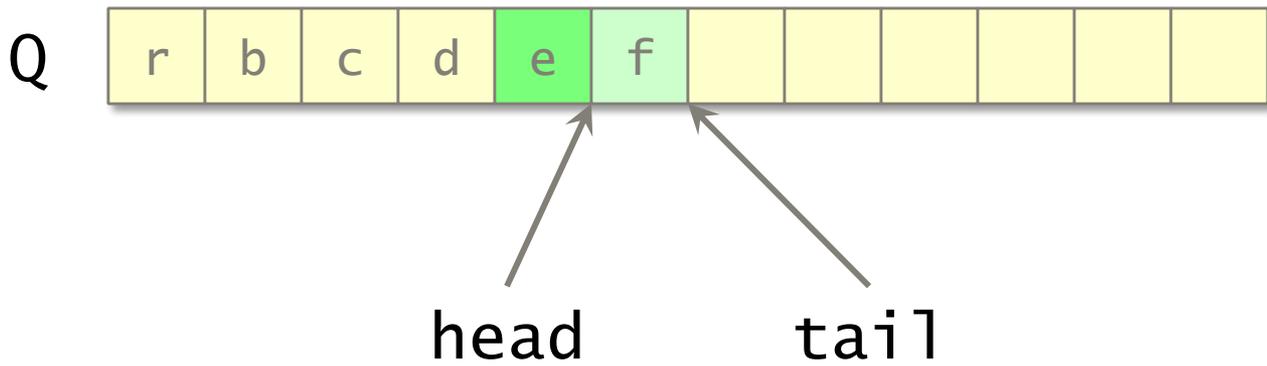
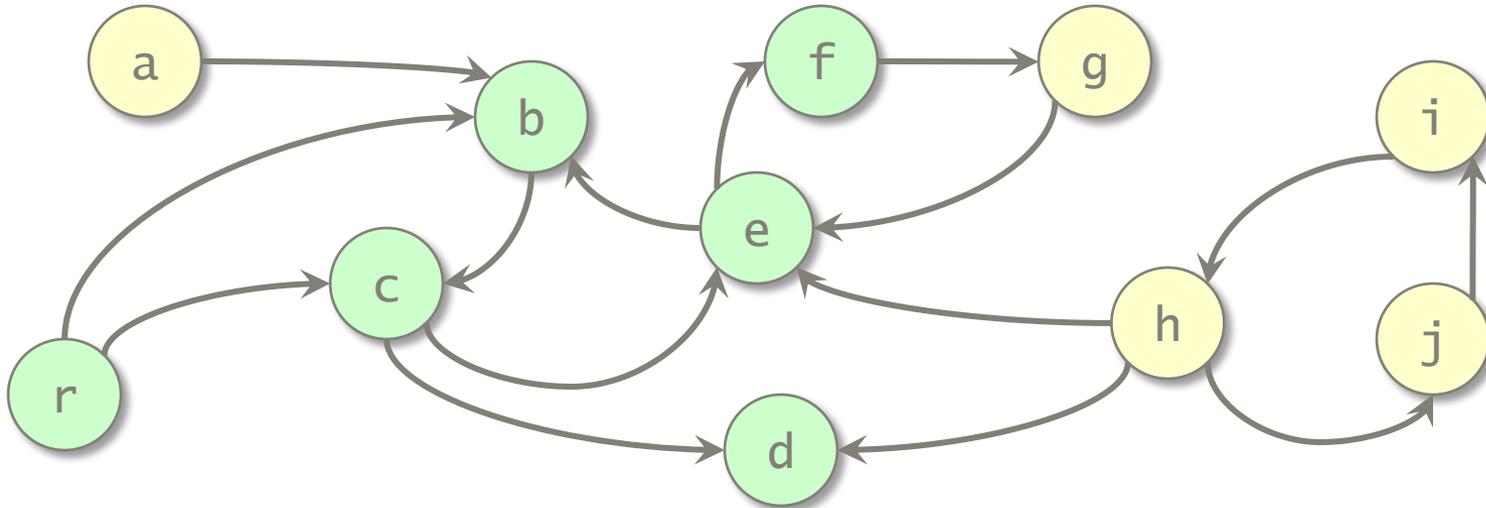
Breadth-First Search



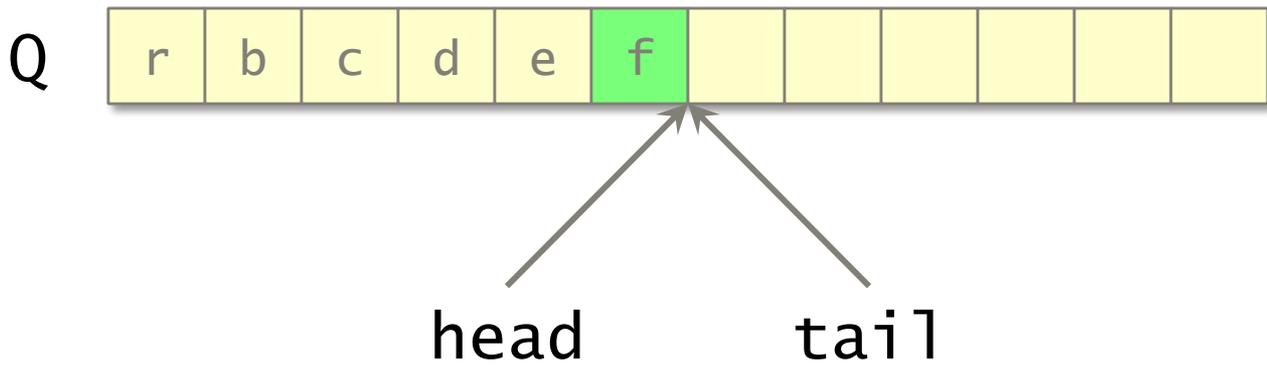
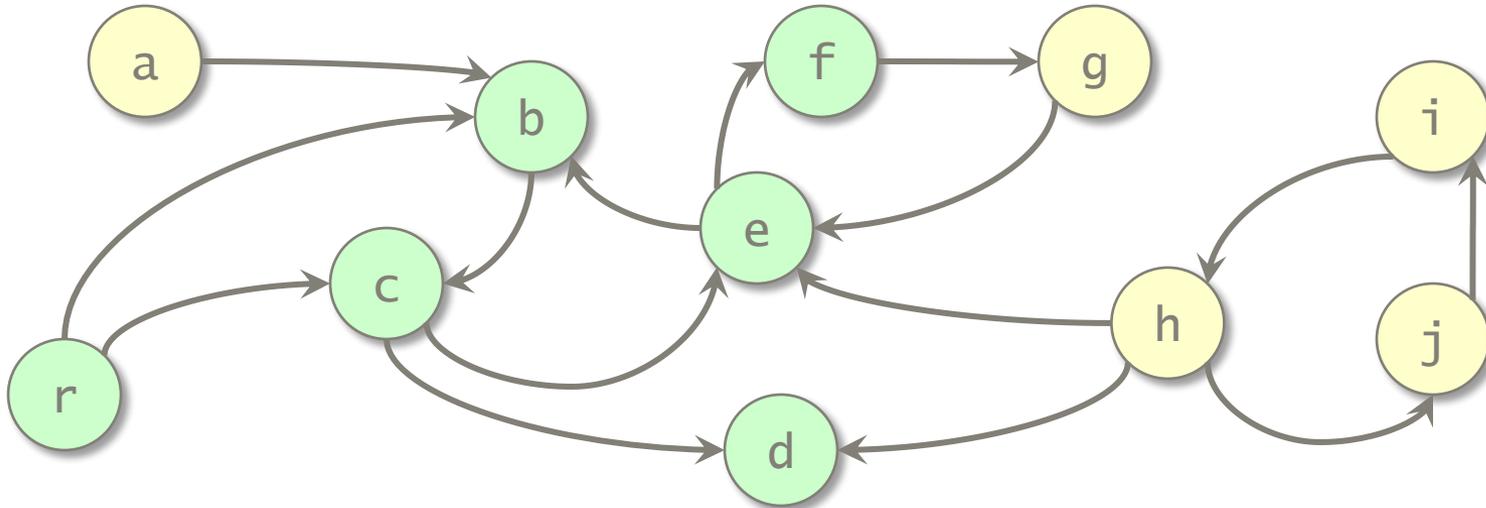
Breadth-First Search



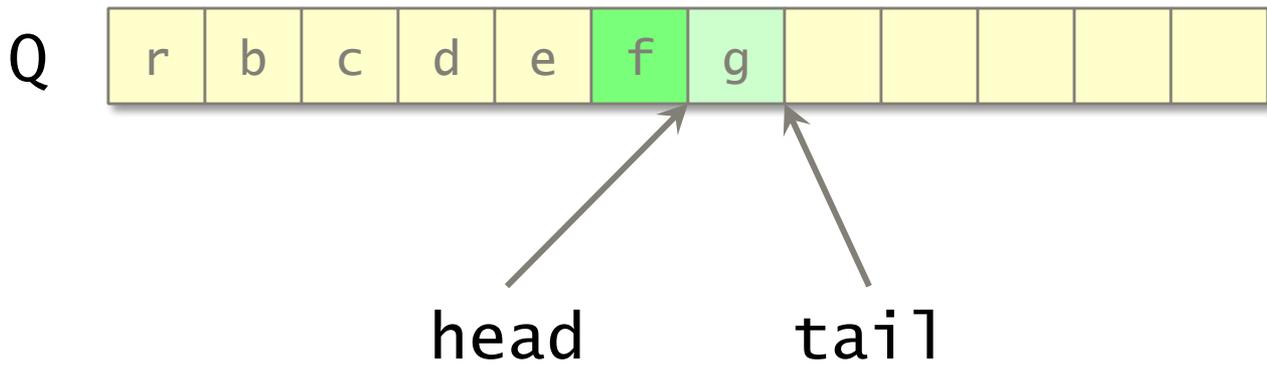
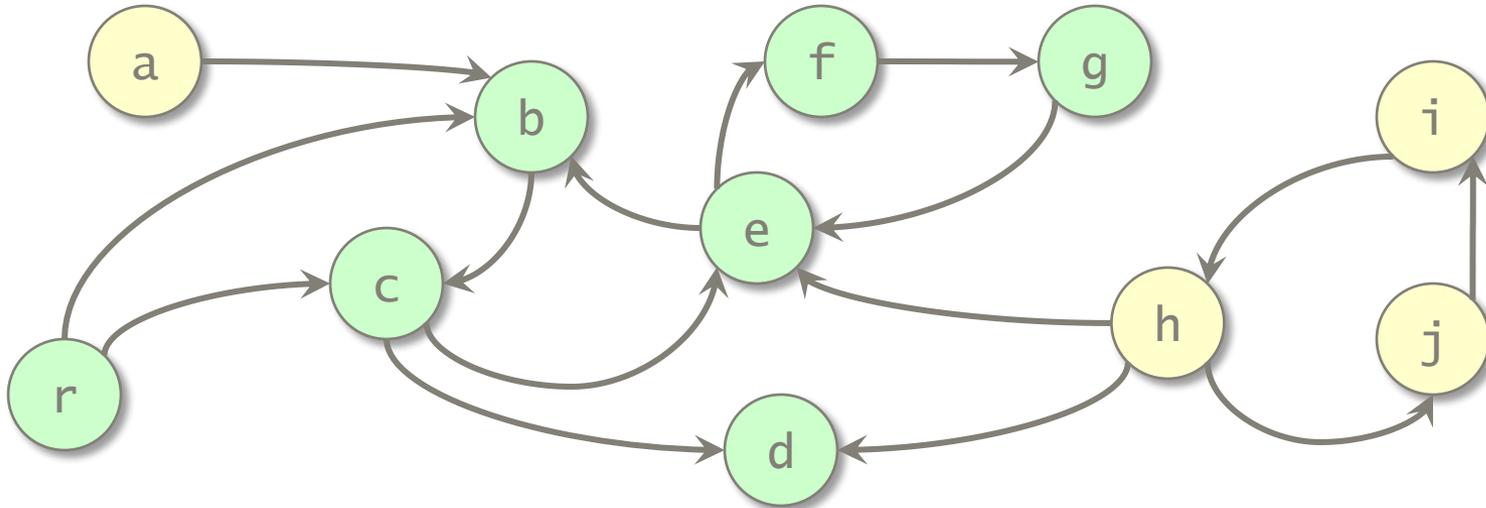
Breadth-First Search



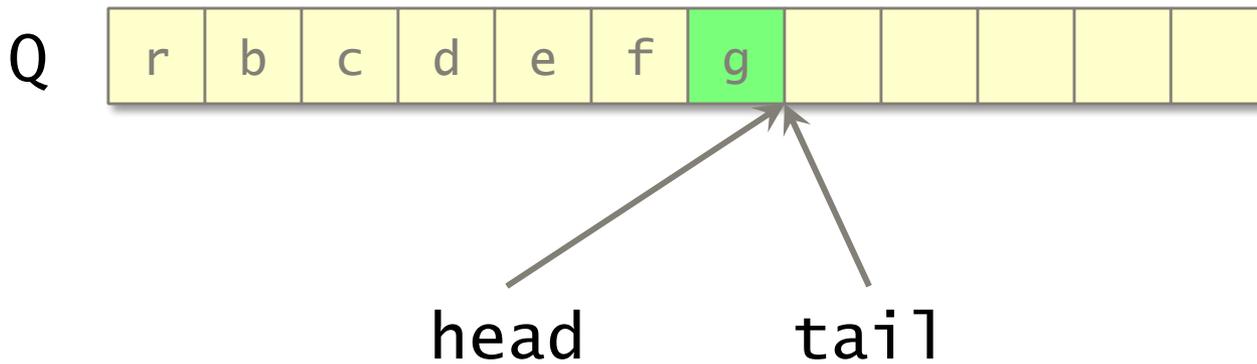
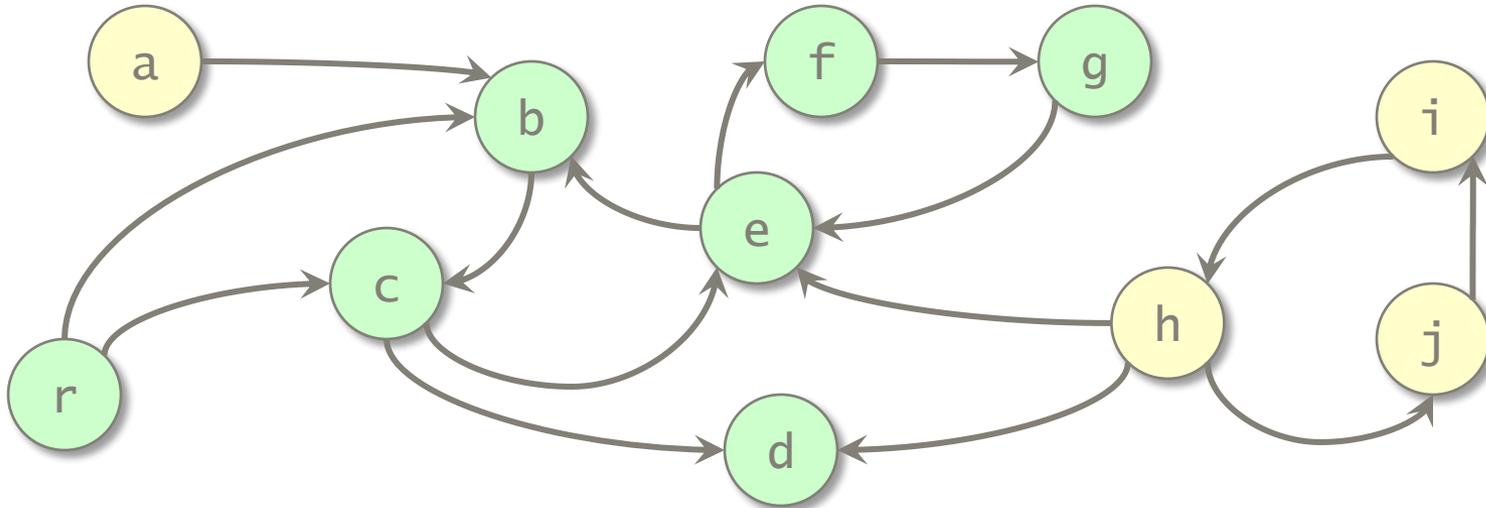
Breadth-First Search



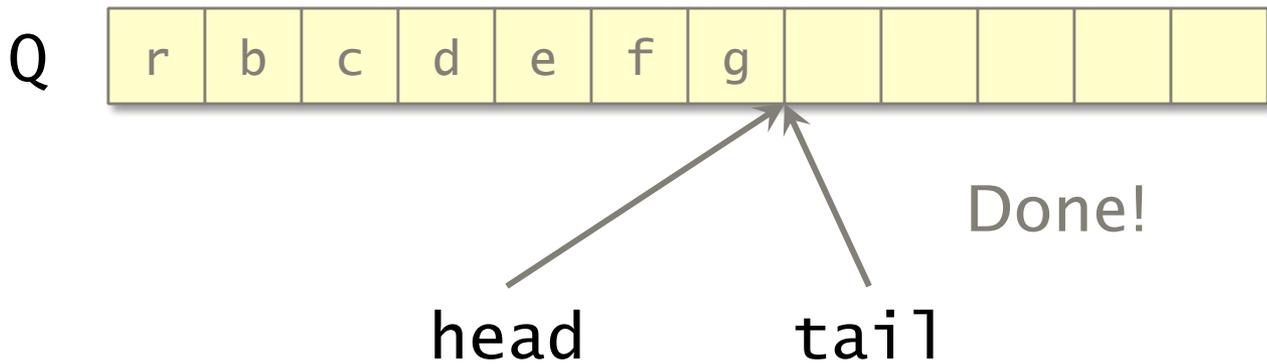
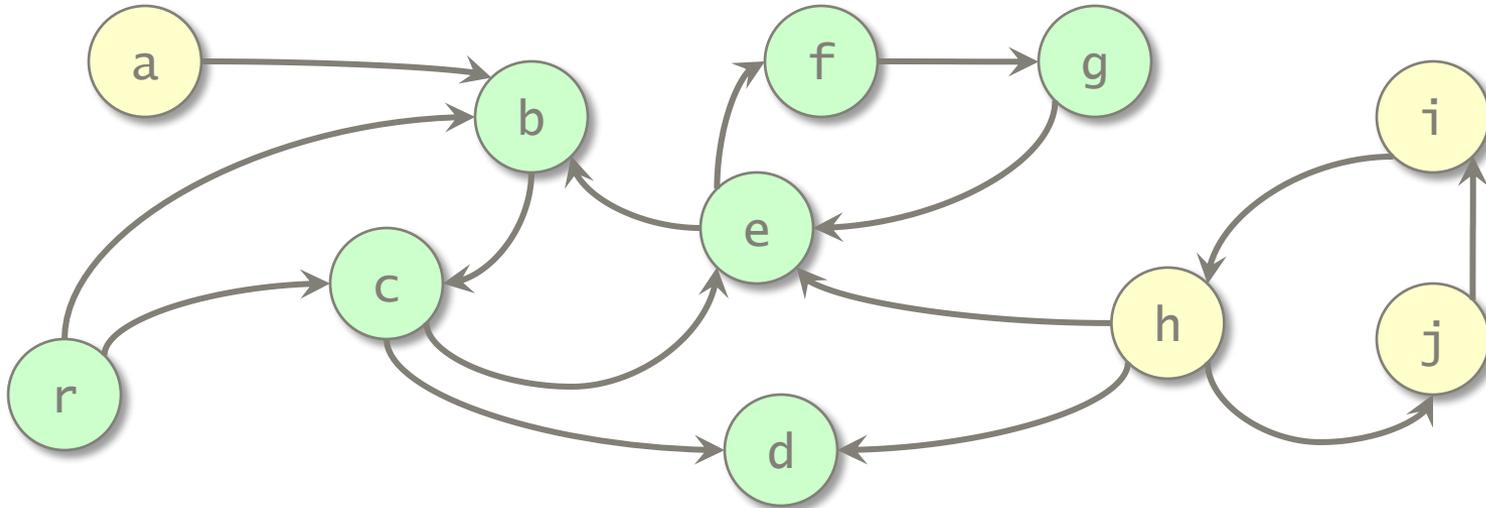
Breadth-First Search



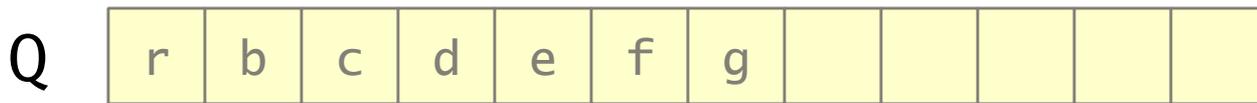
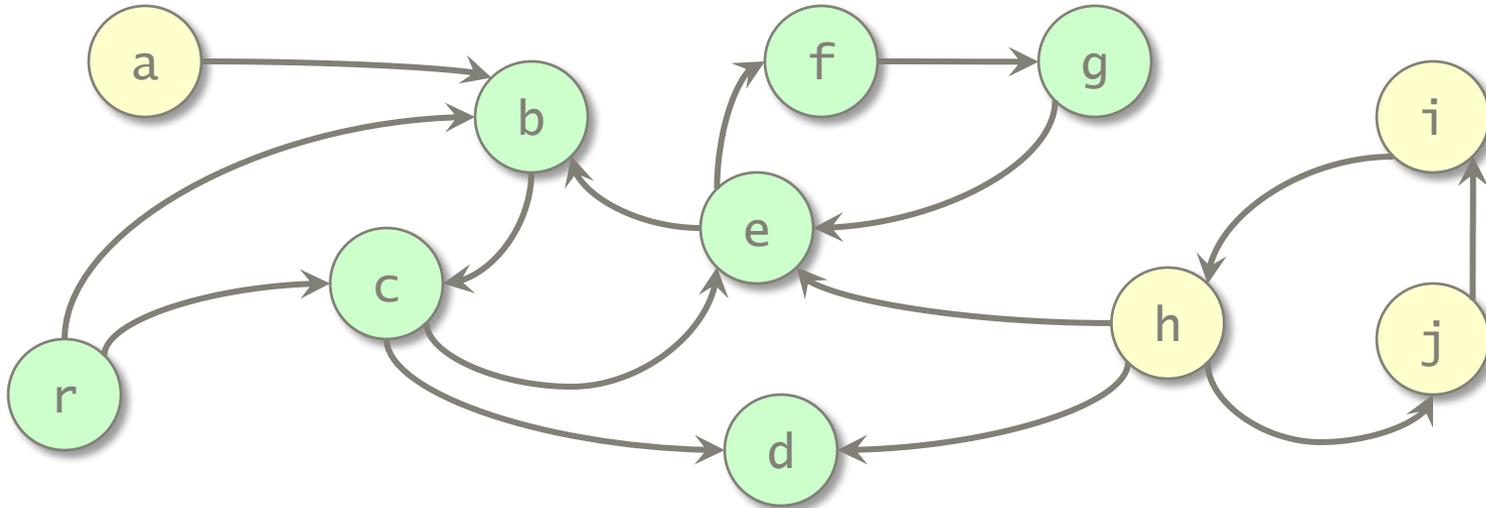
Breadth-First Search



Breadth-First Search



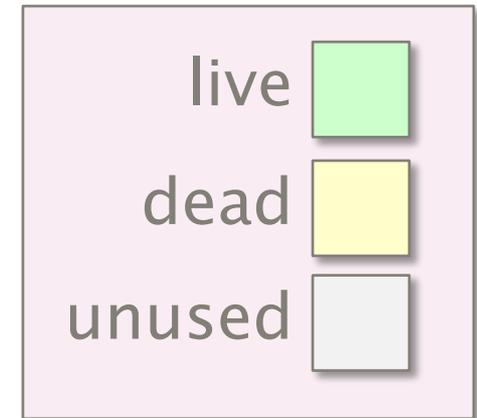
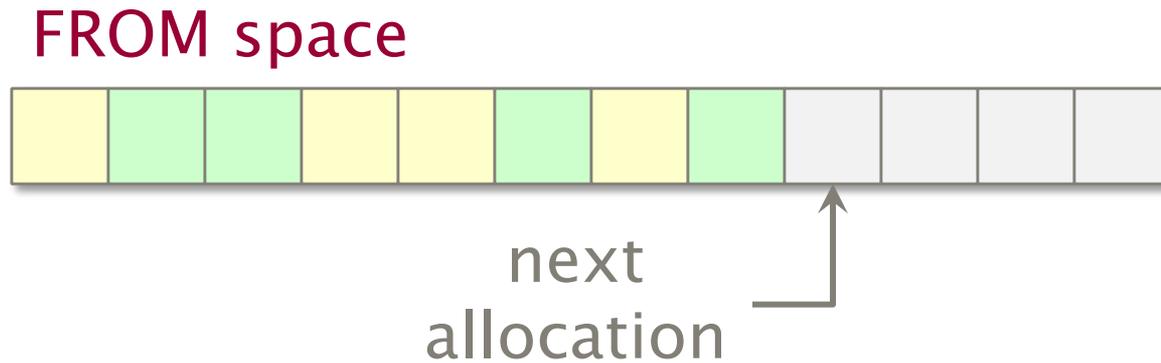
Breadth-First Search



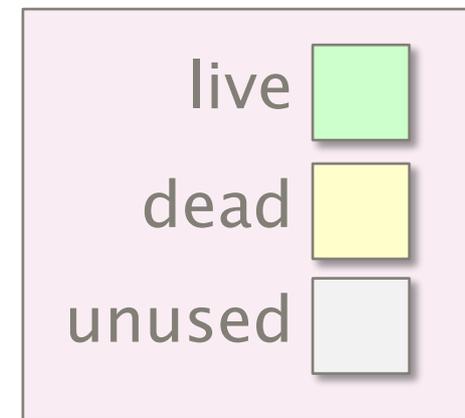
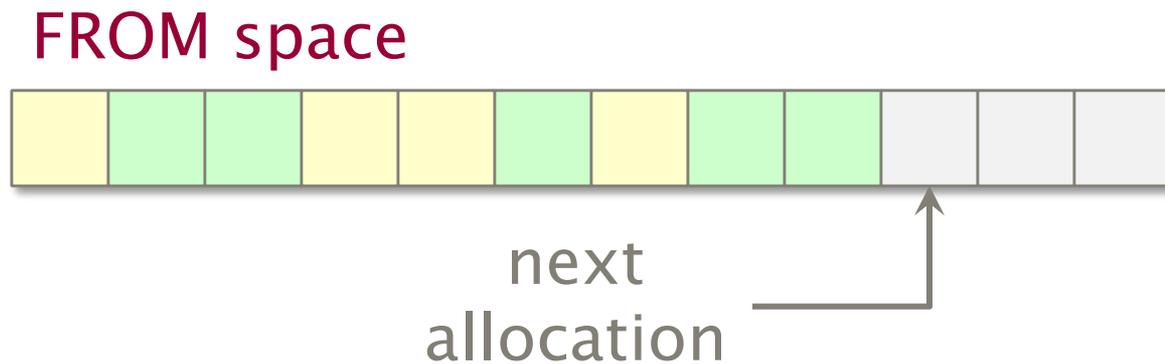
Observation

All live vertices are placed in contiguous storage in Q.

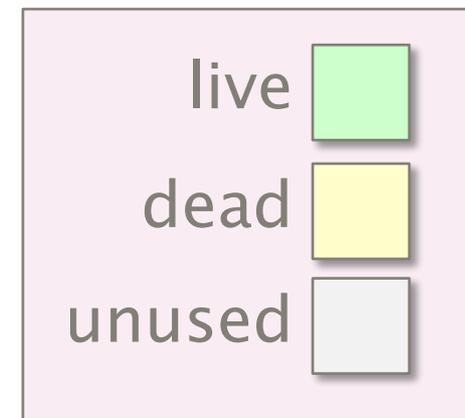
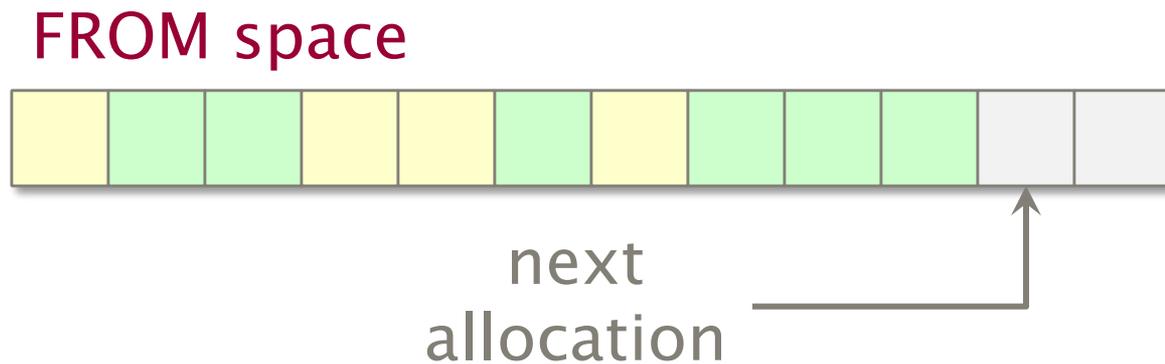
Copying Garbage Collector



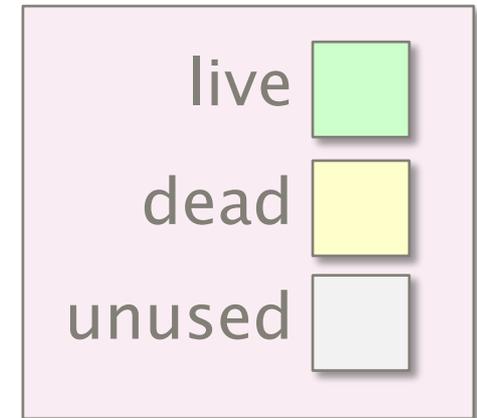
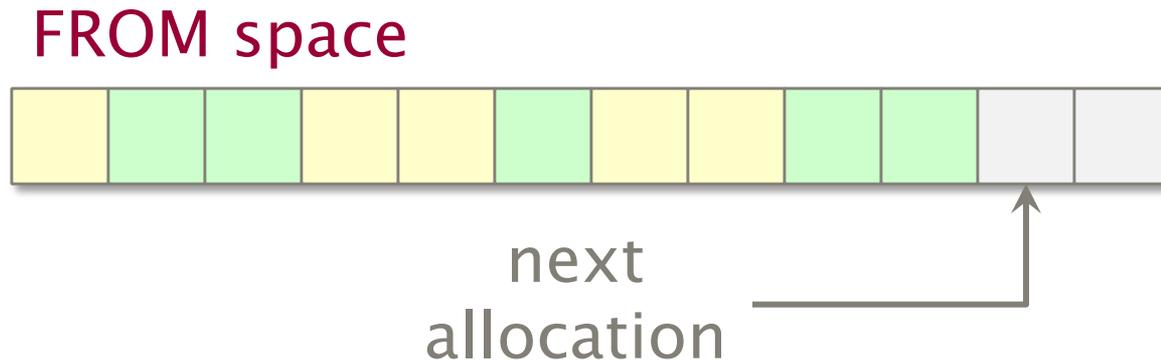
Copying Garbage Collector



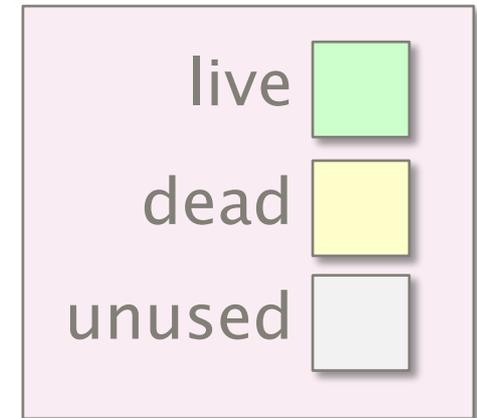
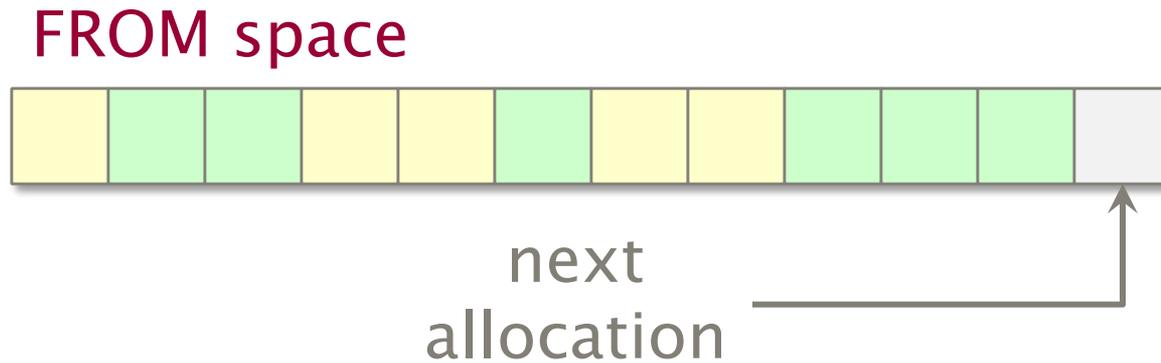
Copying Garbage Collector



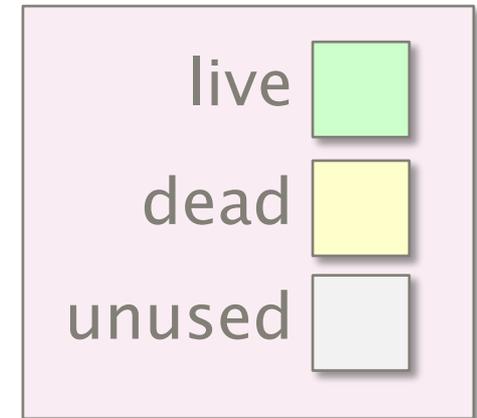
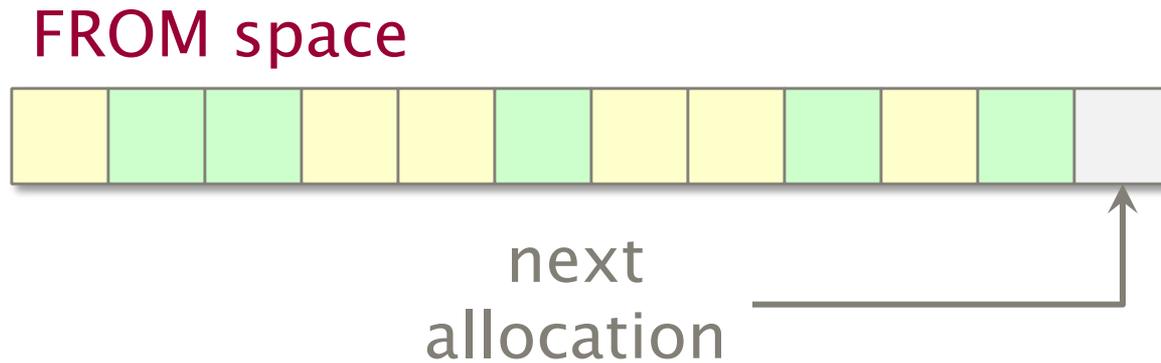
Copying Garbage Collector



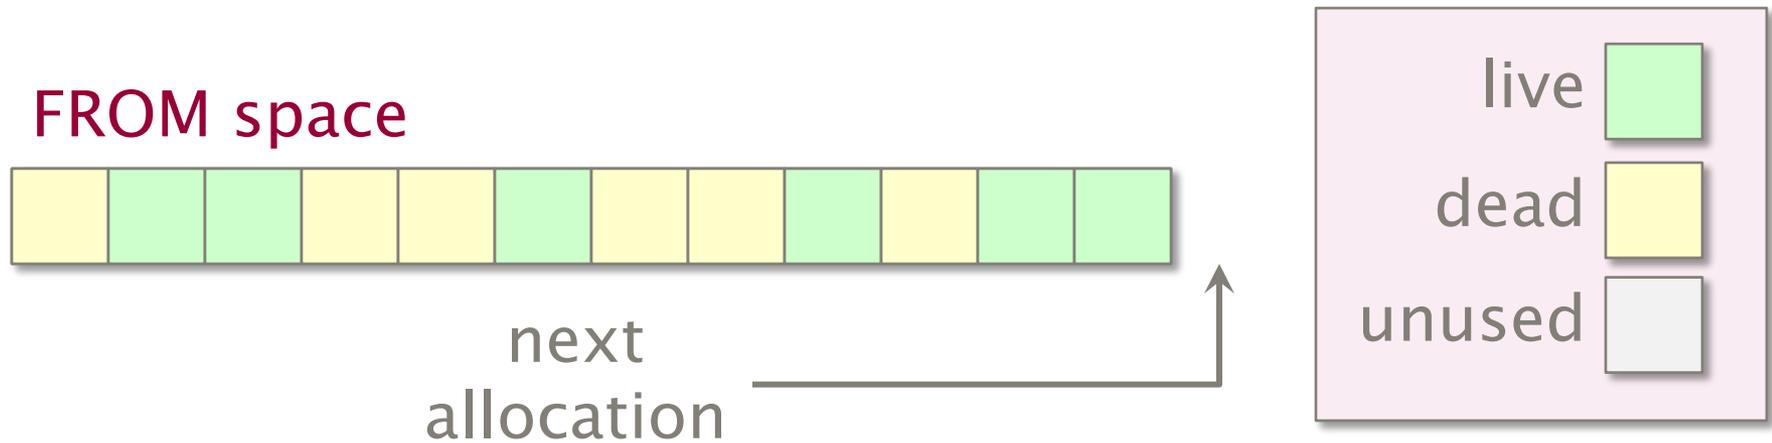
Copying Garbage Collector



Copying Garbage Collector

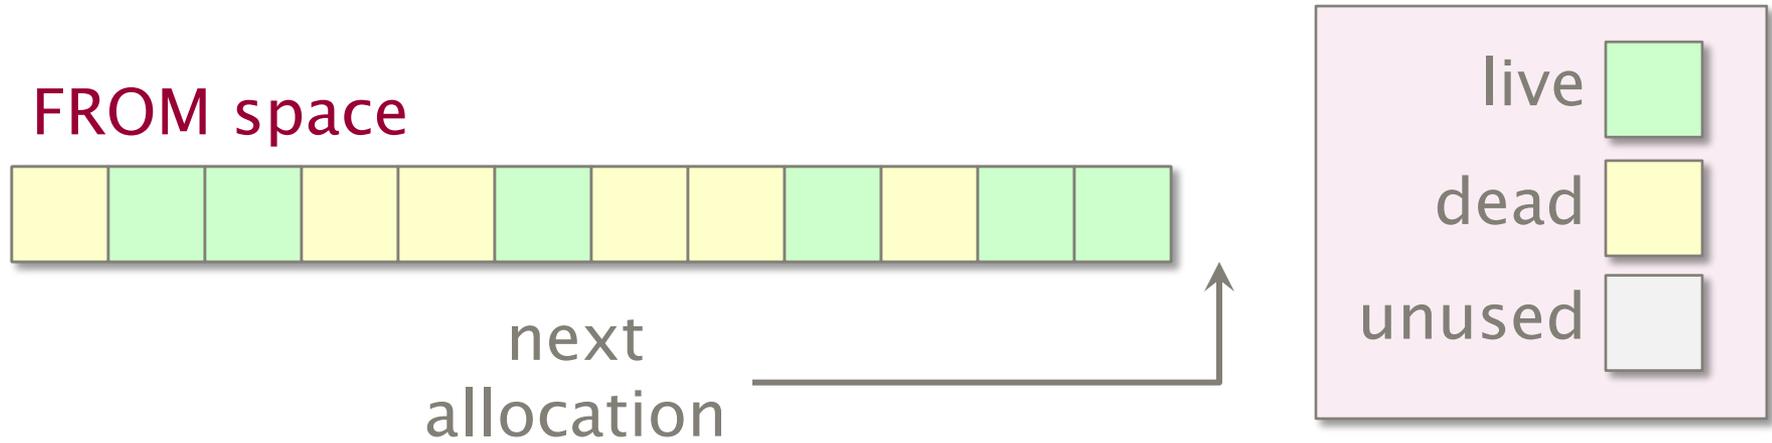


Copying Garbage Collector

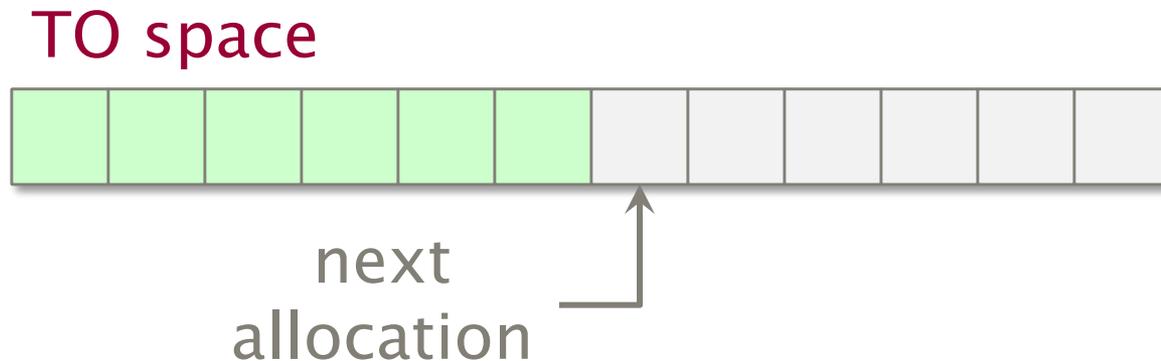


When the FROM space is full, copy live storage using BFS with the TO space as the FIFO queue.

Copying Garbage Collector



When the FROM space is full, copy live storage using BFS with the TO space as the FIFO queue.

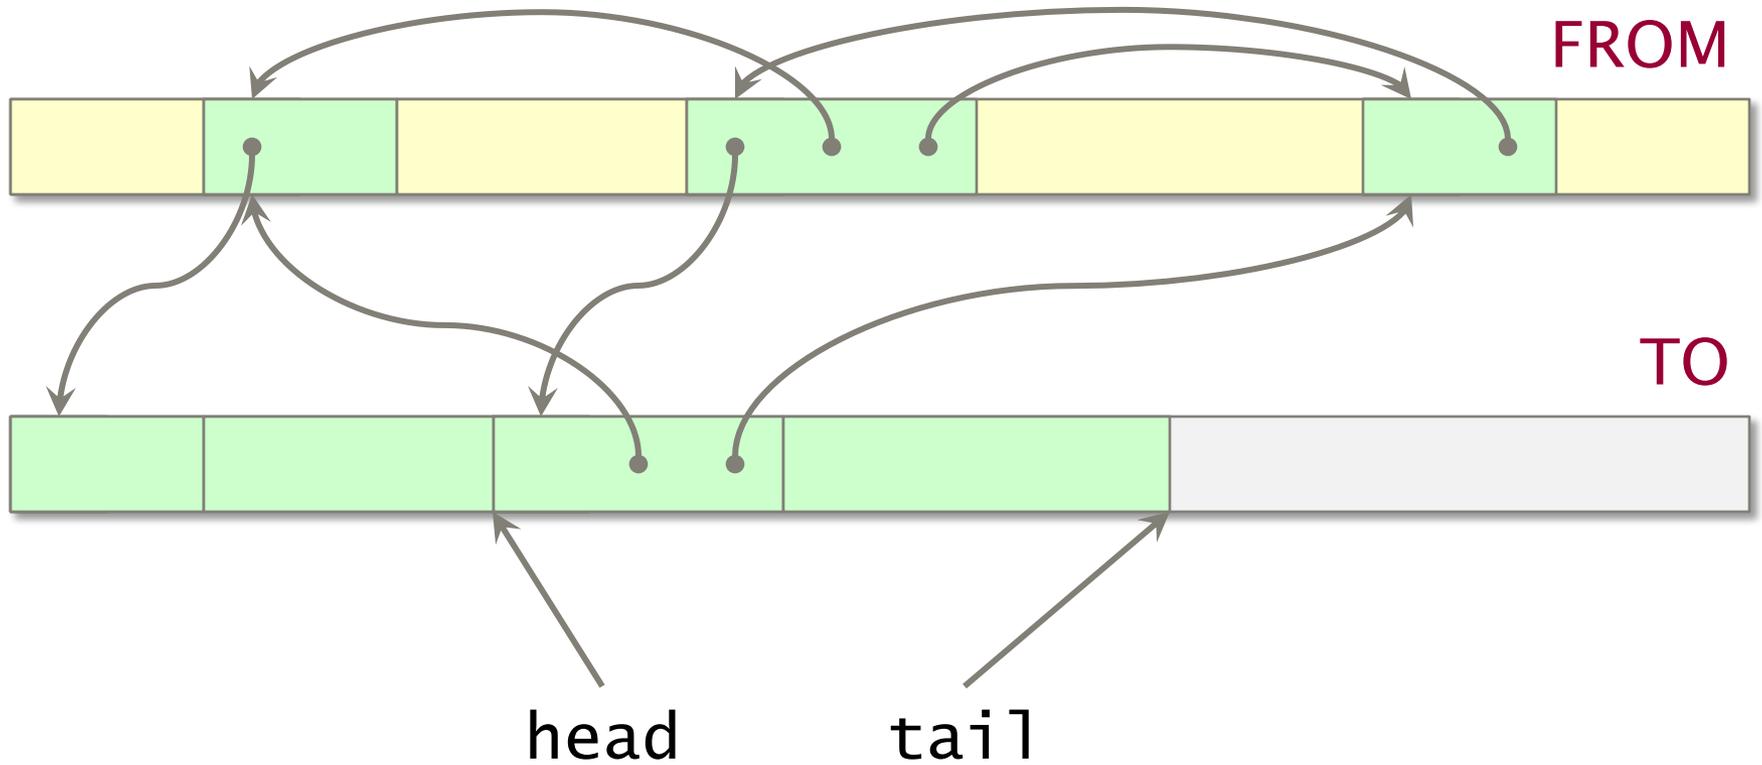


Updating Pointers

Since the FROM address of an object is not generally equal to the TO address of the object, pointers must be updated.

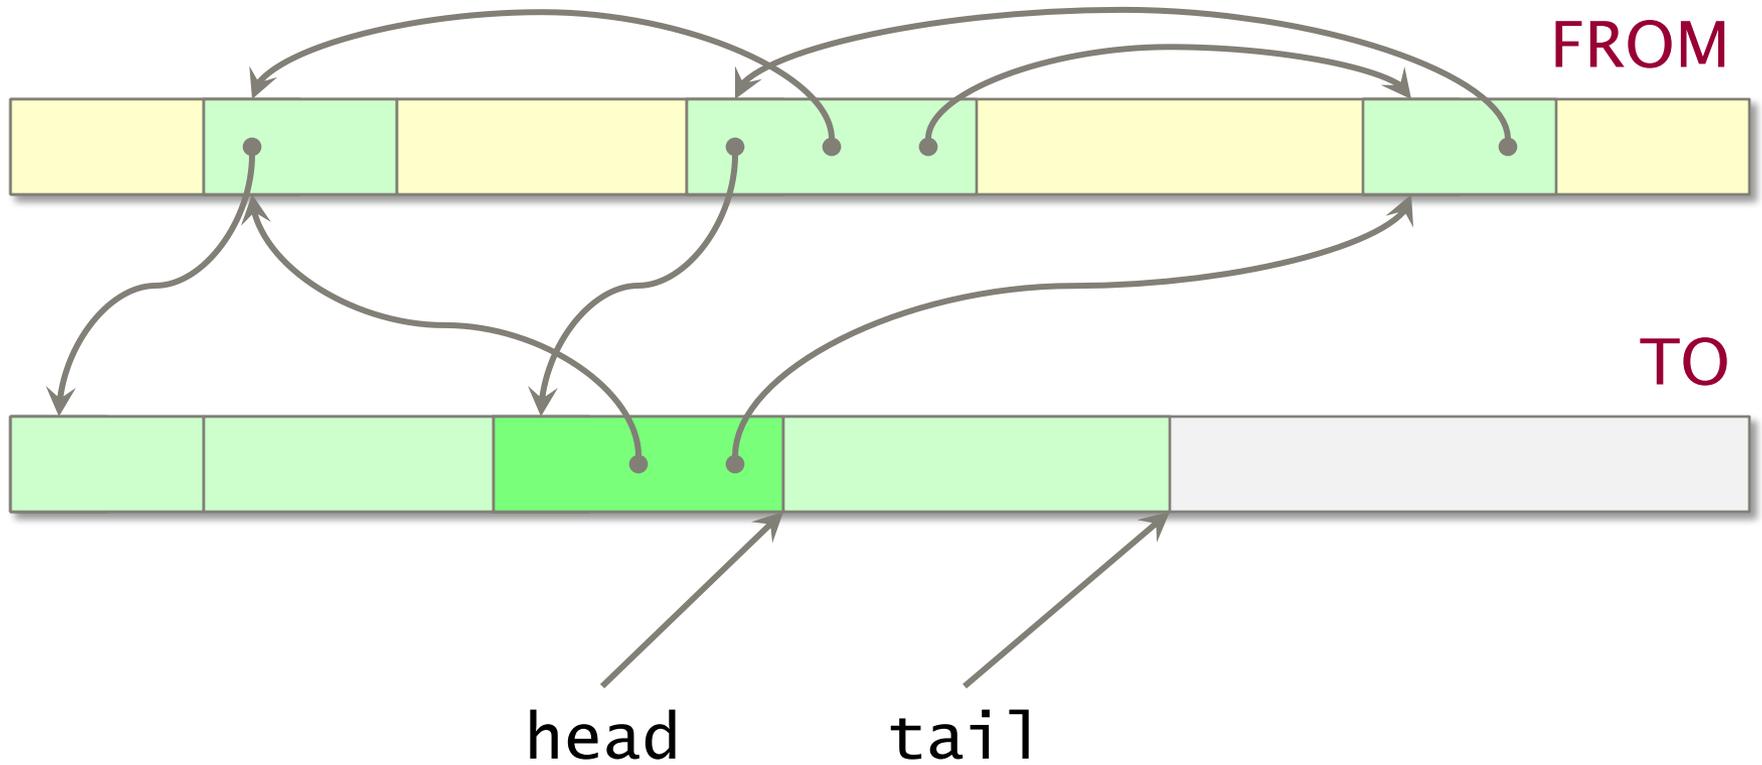
- When an object is copied to the TO space, store a forwarding pointer in the FROM object, which implicitly marks it as moved.
- When an object is removed from the FIFO queue in the TO space, update all its pointers.

Example



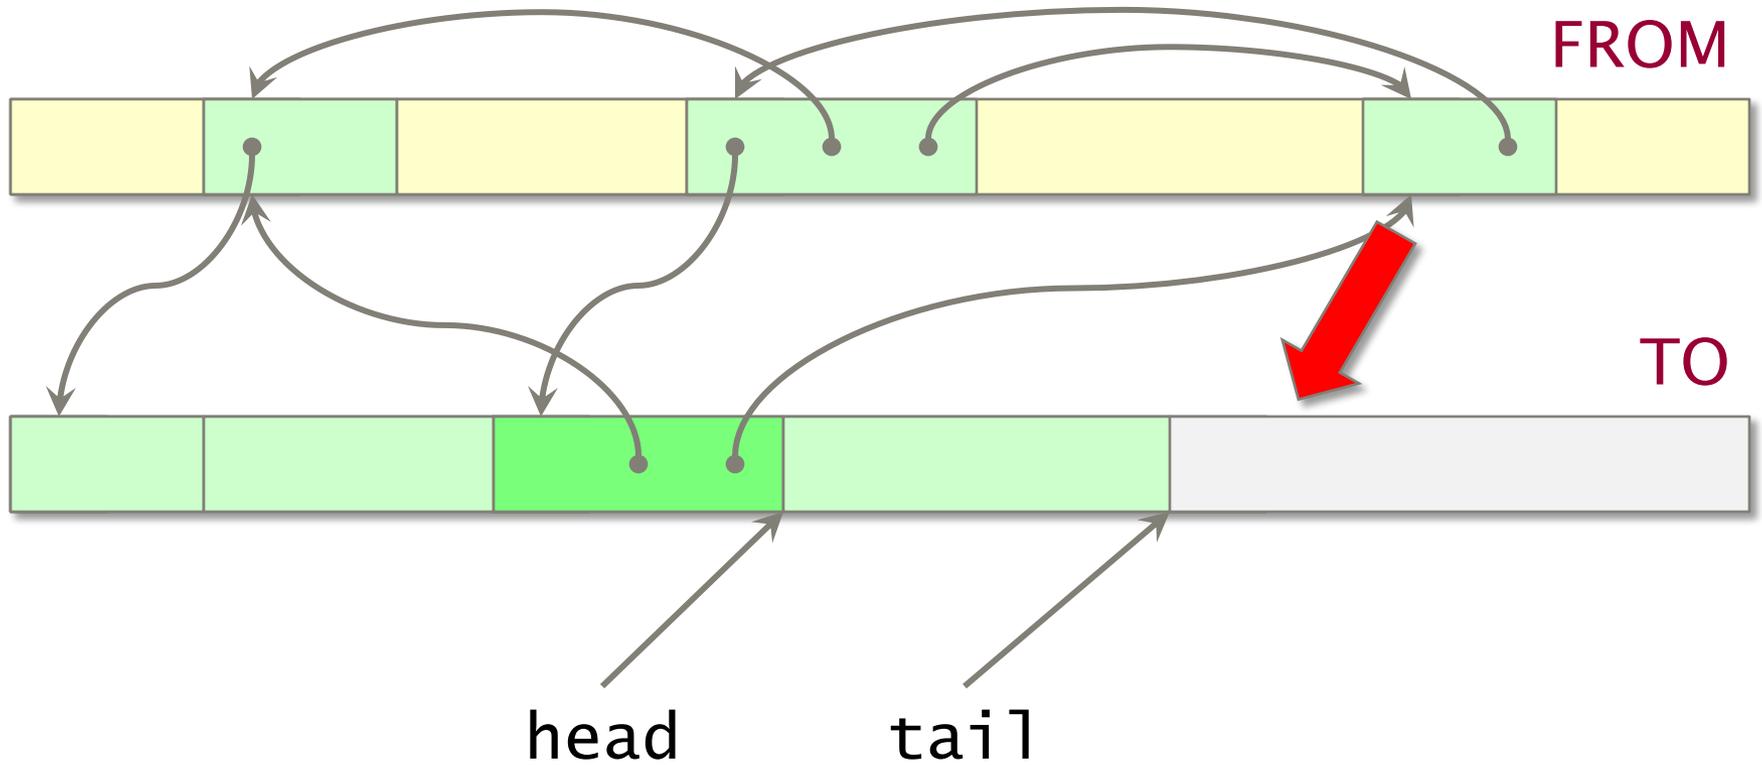
Remove an item from the queue.

Example



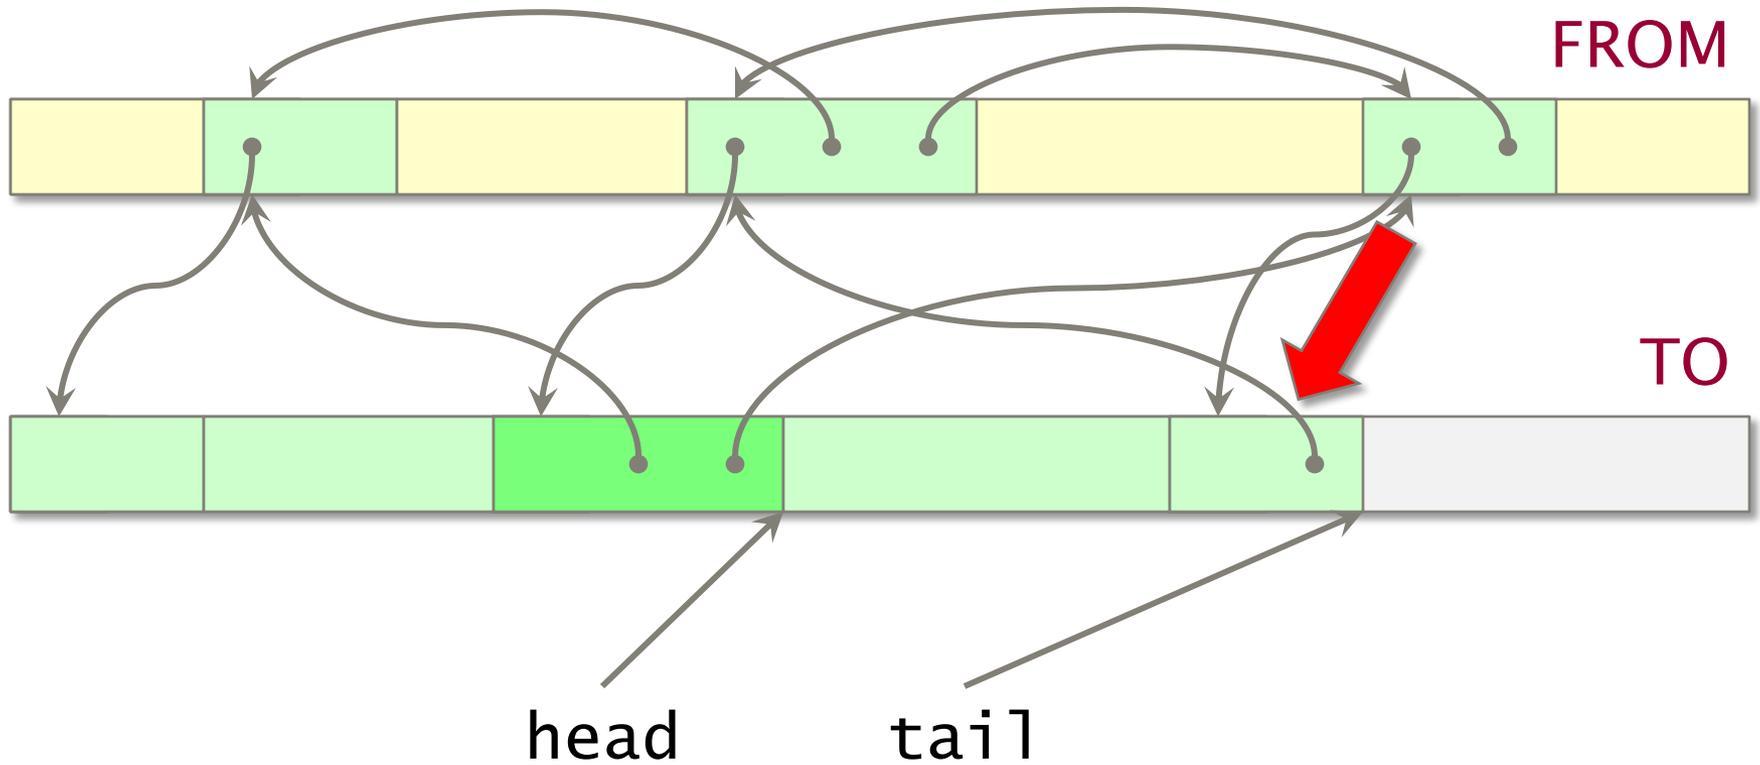
Remove an item from the queue.

Example



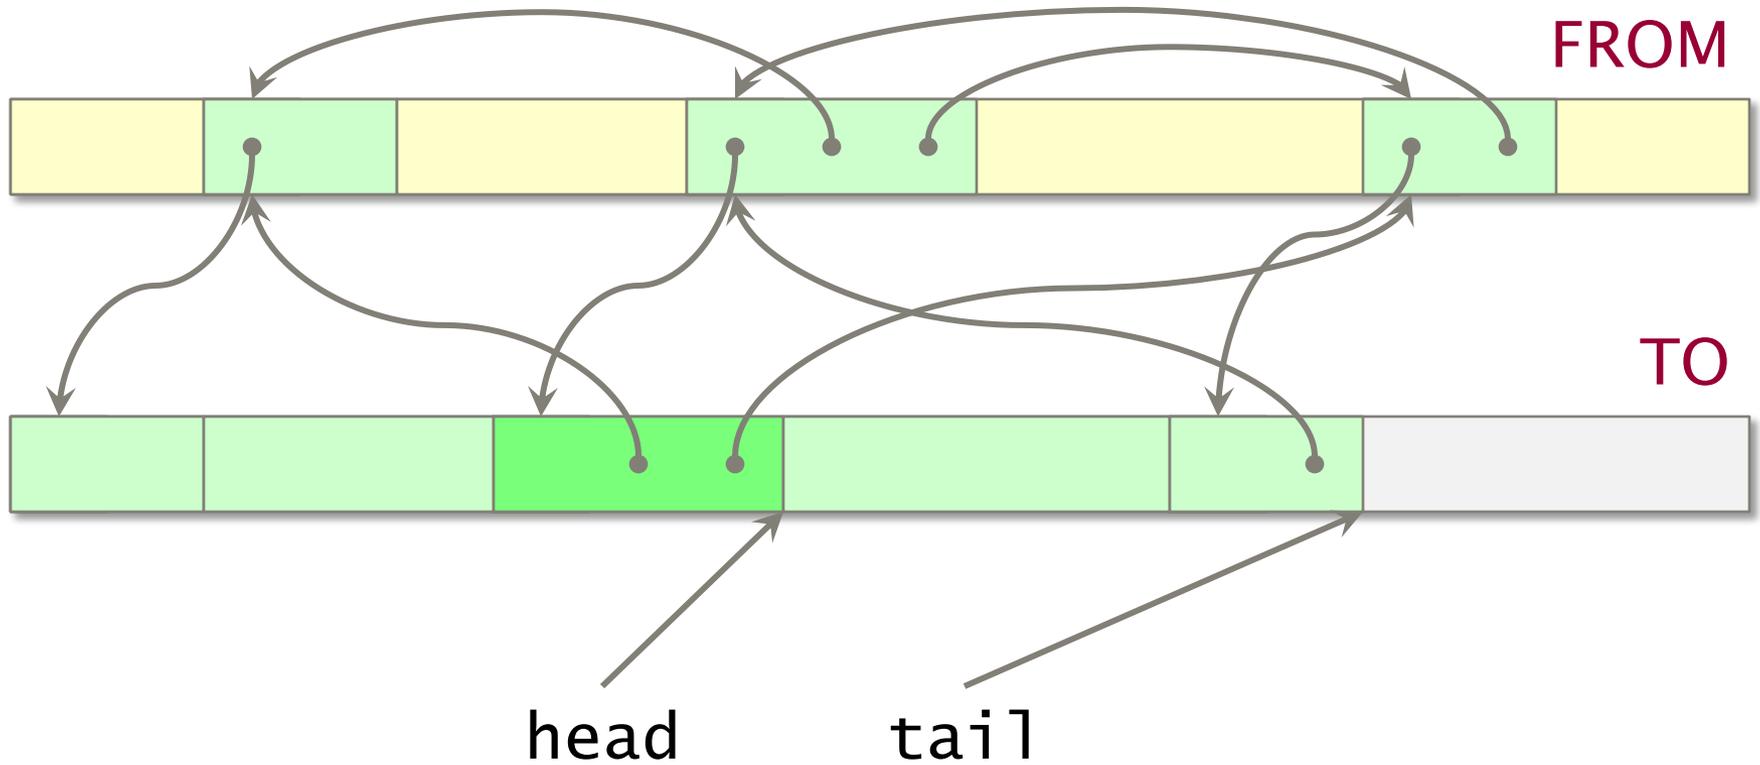
Enqueue adjacent vertices.

Example



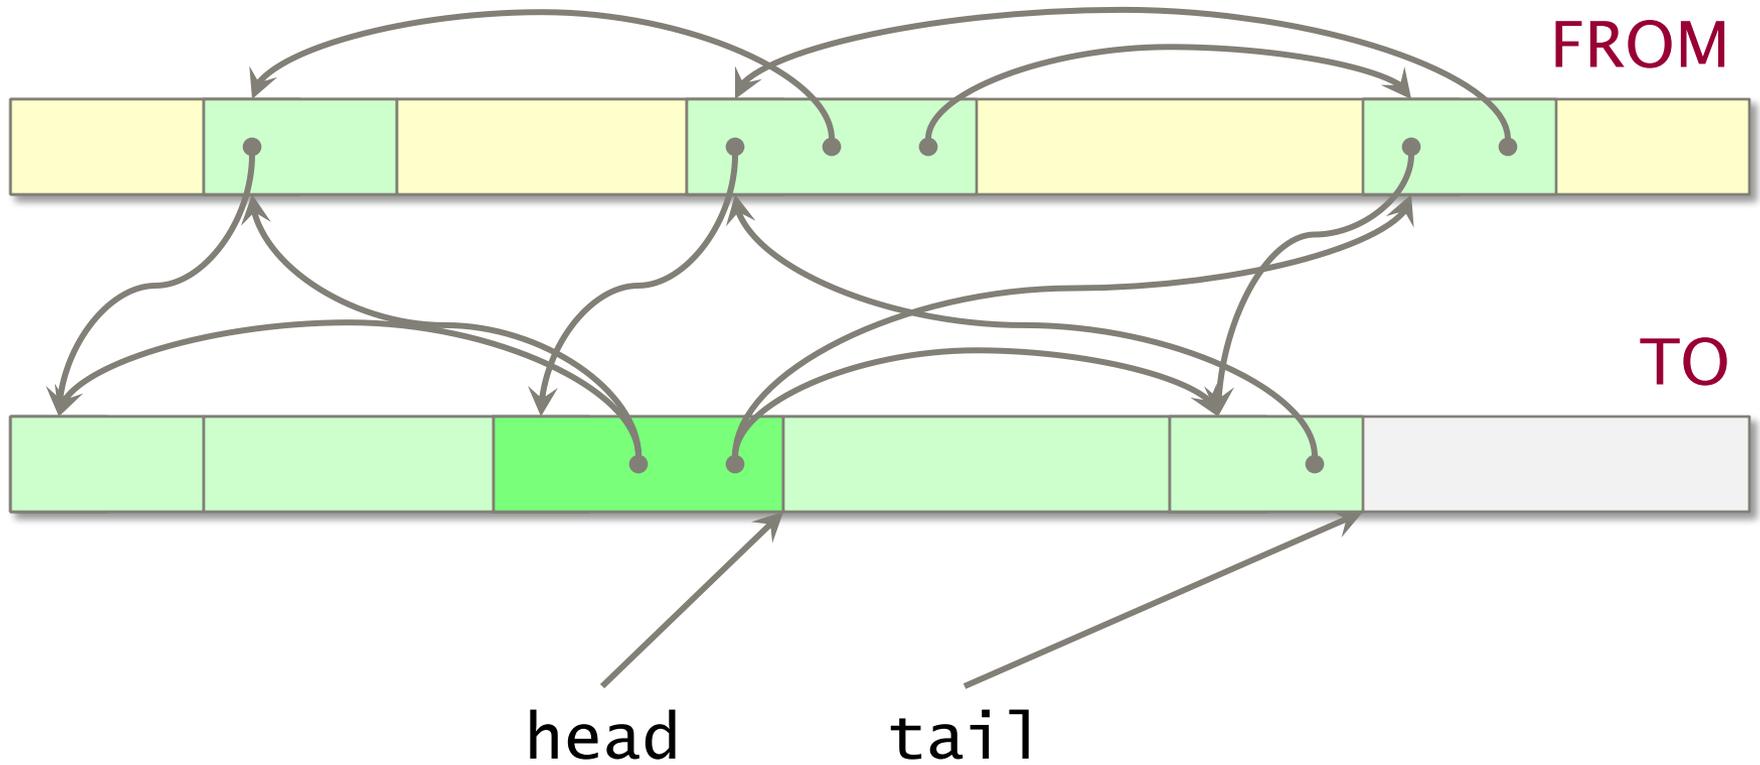
Enqueue adjacent vertices.
Place forwarding pointers in FROM vertices.

Example



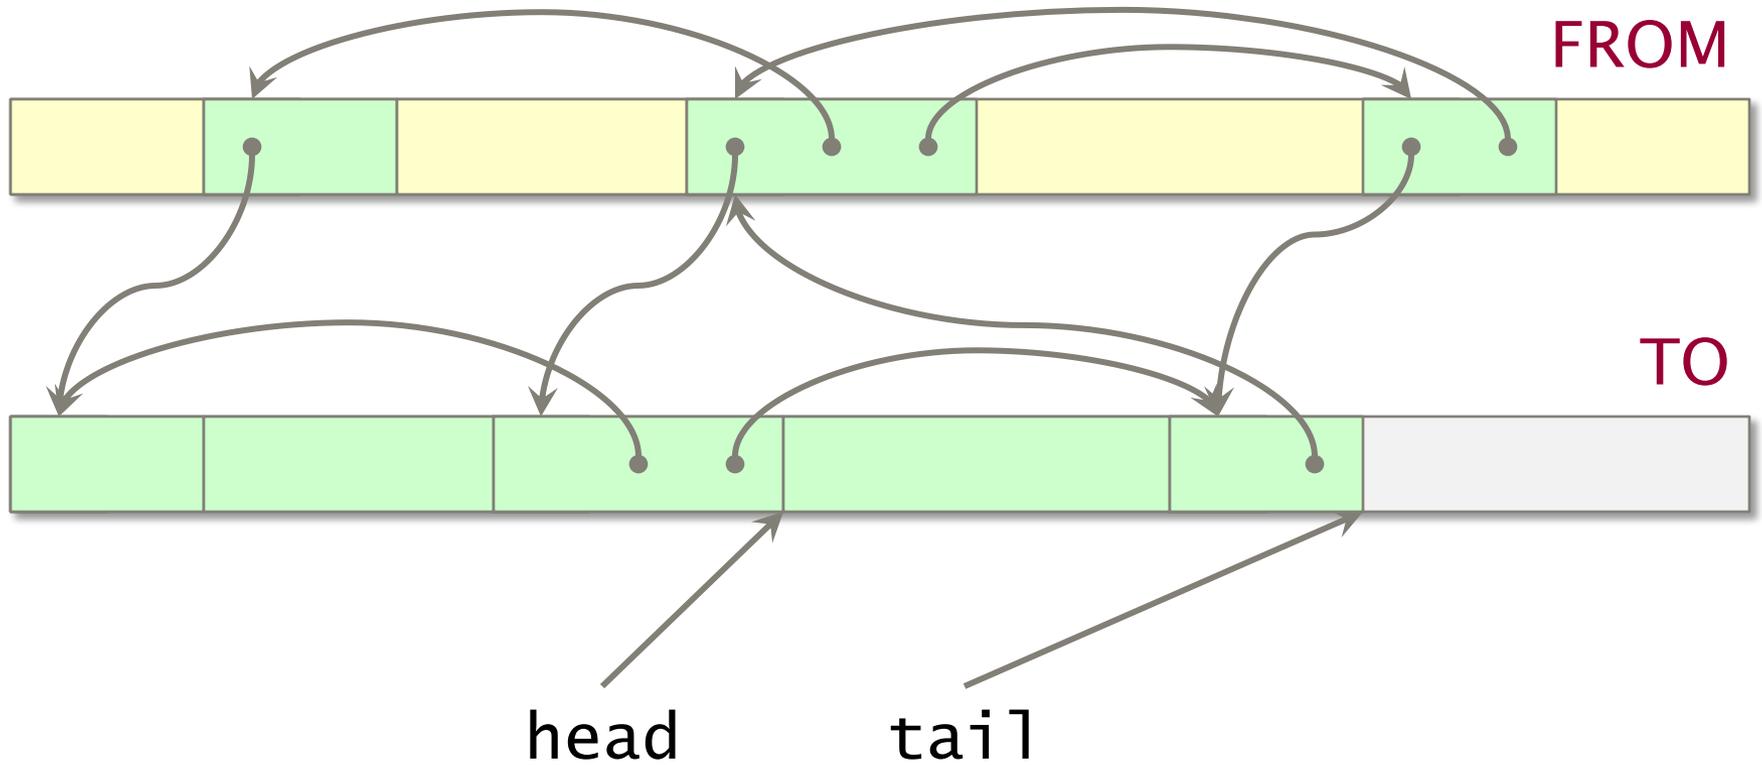
Update the pointers in the removed item to refer to its adjacent items in the TO space.

Example



Update the pointers in the removed item to refer to its adjacent items in the TO space.

Example



Linear time to copy and update all vertices.

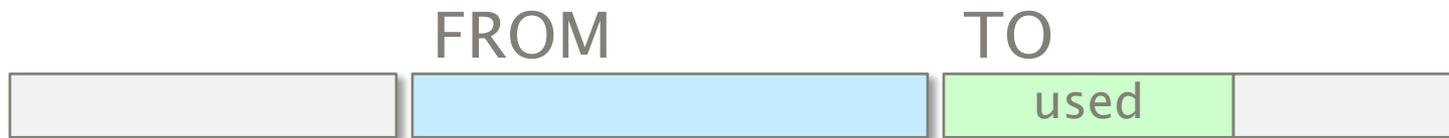
Managing Virtual Memory



After copying, the old TO becomes the new FROM. Allocate new heap space equal to the used space in the new FROM.



When space runs out, the new TO is allocated with the same size as FROM. Allocate TO at start if it fits. Otherwise, allocate TO after FROM.



Theorem. VM space used is $\Theta(1)$ times optimal. ■

Dynamic Storage Allocation

Lots more is known and unknown about dynamic storage allocation. Strategies include

- buddy system,
- mark-and-sweep garbage collection,
- generational garbage collection,
- real-time garbage collection,
- multithreaded storage allocation,
- parallel garbage collection,
- etc.

MIT OpenCourseWare
<http://ocw.mit.edu>

6.172 Performance Engineering of Software Systems
Fall 2010

For information about citing these materials or our Terms of Use, visit: <http://ocw.mit.edu/terms>.