

floordrobe, *noun*. A heap of clothing left on the floor of a room.
In computer science: the most perfect design for an advanced data structure.



Pushing N items is $O(N \log N)$ — but if we're clever we can create a binary heap of N items in $O(N)$.

	popmin	push	decreasekey
binary heap	$O(\log N)$	$O(\log N)$	$O(\log N)$
binomial heap	$O(\log N)$	$O(\log N)$ $O(1)$ amortized	$O(\log N)$

Dijkstra's algorithm makes $O(E)$ calls to push / decreasekey, and only $O(V)$ calls to popmin.

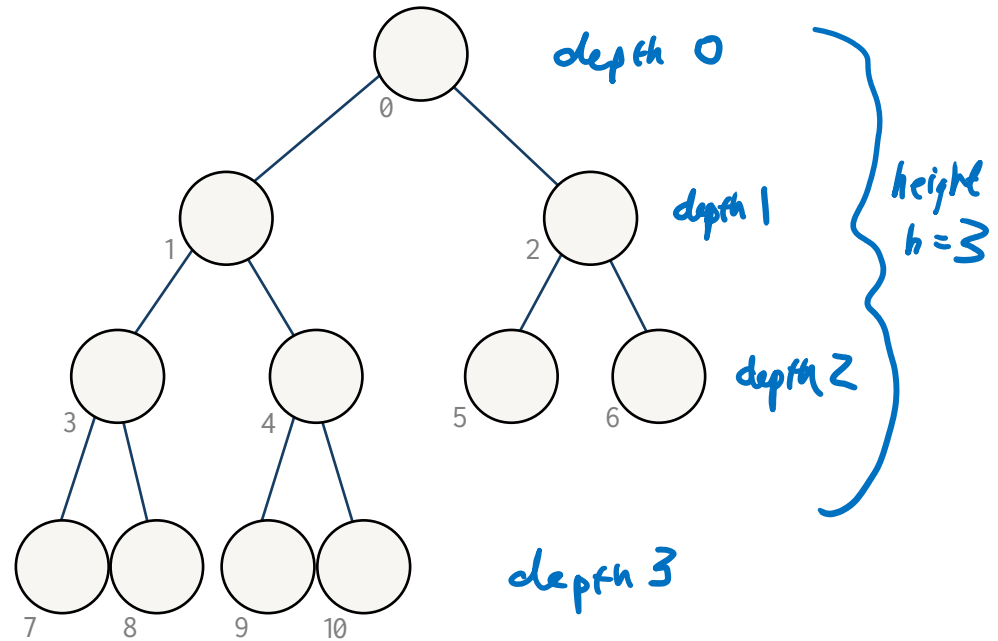
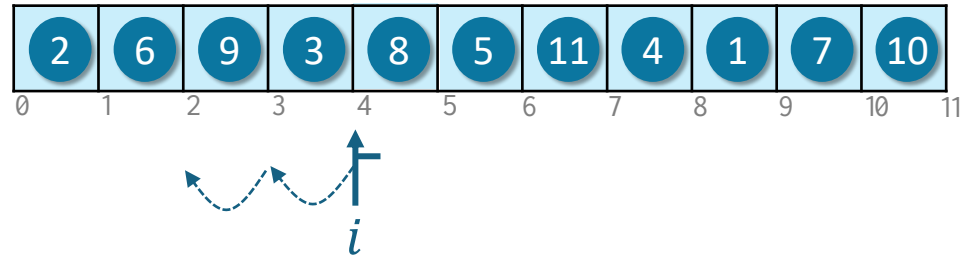
QUESTION1. Can we make both push and decreasekey be $O(1)$?

QUESTION2. What's the binomial heap's secret sauce that lets it have $O(1)$ push?

Pushing N items is $O(N \log N)$ — but if we're clever we can create a binary heap of N items in $O(N)$.

	popmin	push	decreasekey
binary heap	$O(\log N)$	$O(\log N)$	$O(\log N)$
binomial heap	$O(\log N)$	$O(\log N)$ $O(1)$ amortized	$O(\log N)$

```
# Fast binary-heapification
for i in ([N/2]-1)..0:
    # assert: trees rooted at (i+1)..N are heaps
    re-heapify the tree rooted at x[i]
    by bubbling down
```



- When we reheapify from depth d it takes $h - d$ work to bubble down, and there are $\leq 2^d$ items that need this work.
- There are more items at greater depths, and it's these items that take the least work.
- Total work is $\sum_{d=0}^h 2^d (h - d) \leq 2 \times 2^h = 2N$ [printed notes chapter 2.10]

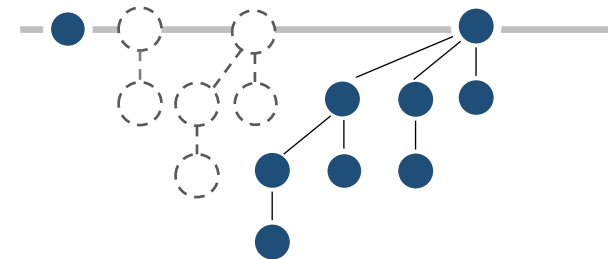
Pushing N items is $O(N \log N)$ — but if we're clever we can create a binary heap of N items in $O(N)$.

	popmin	push	decreasekey
binary heap	$O(\log N)$	$O(\log N)$	$O(\log N)$
binomial heap	$O(\log N)$	$O(\log N)$ $O(1)$ amortized	$O(\log N)$



SECRET SAUCE. Design your data structure so that most of the time it's sufficient to only touch a small bit of it.

- The binary heap's fast-heapification achieves this through doing its work in a batch (rather than push by push)
- The binomial heap achieves this by splitting up the heap into semi-isolatable trees



Pushing N items is $O(N \log N)$ — but if we're clever we can create a binary heap of N items in $O(N)$.

	popmin	push	decreasekey
binary heap	$O(\log N)$	$O(\log N)$	$O(\log N)$
binomial heap	$O(\log N)$	$O(\log N)$ $O(1)$ amortized	$O(\log N)$

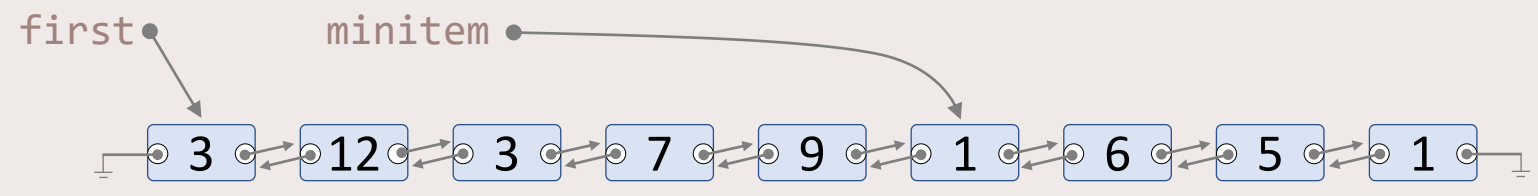
Dijkstra's algorithm makes $O(E)$ calls to push / decreasekey, and only $O(V)$ calls to popmin.

QUESTION1. Can we make both push and decreasekey be $O(1)$?

QUESTION2. What's the binomial heap's secret sauce that lets it have $O(1)$ push?

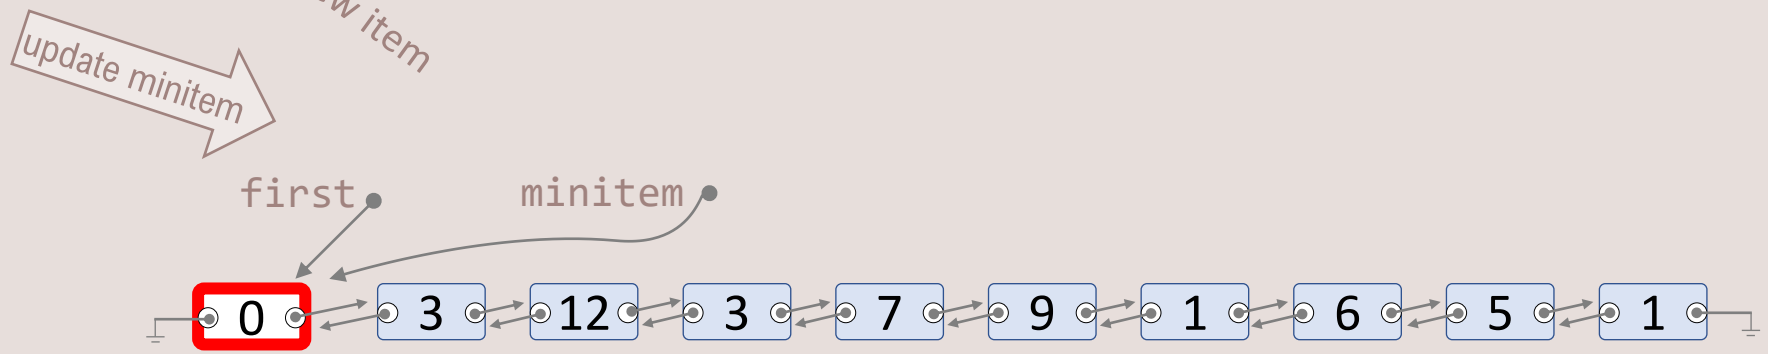
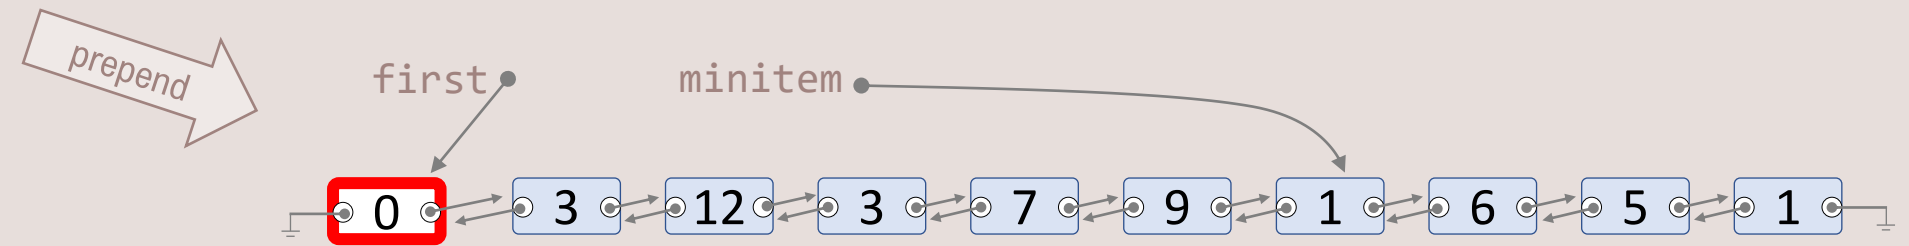
Floordrobe.

Linked-list priority queue

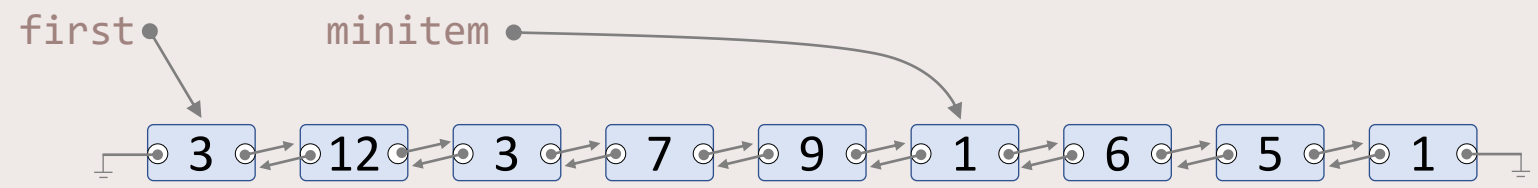


push is $O(1)$

push(*new item*)

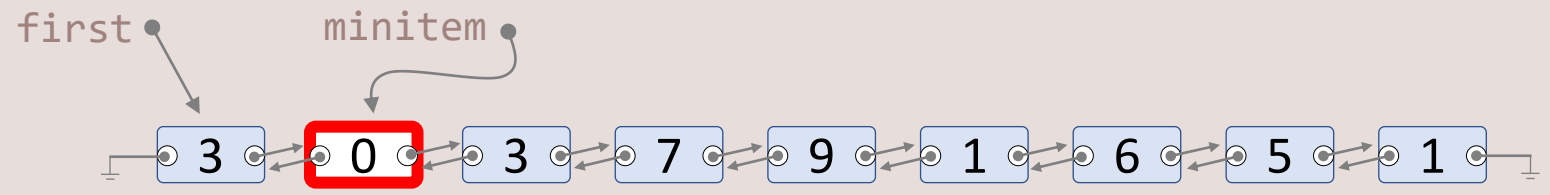
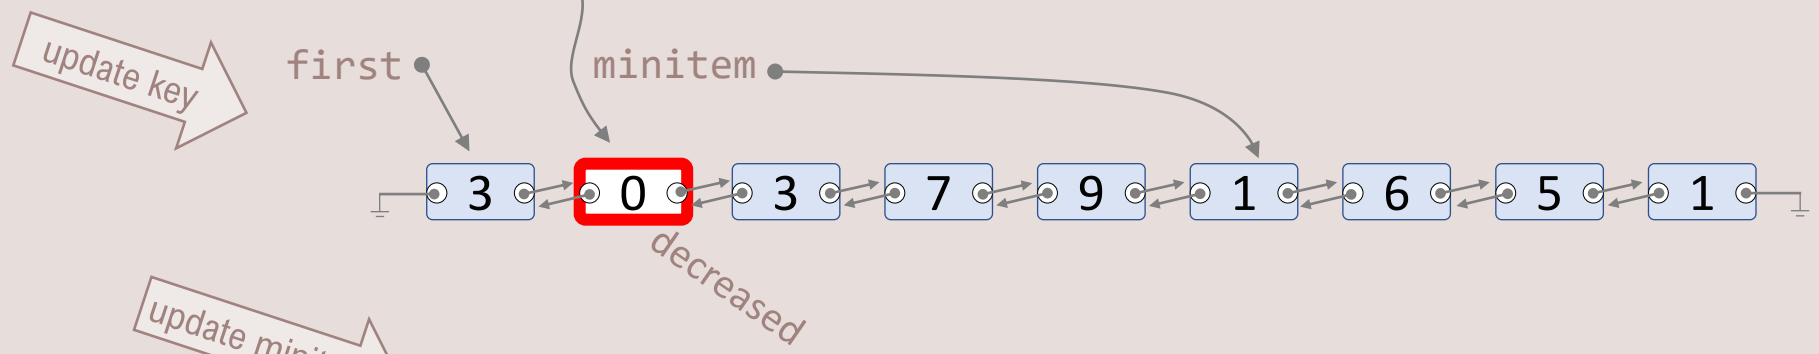


Linked-list priority queue

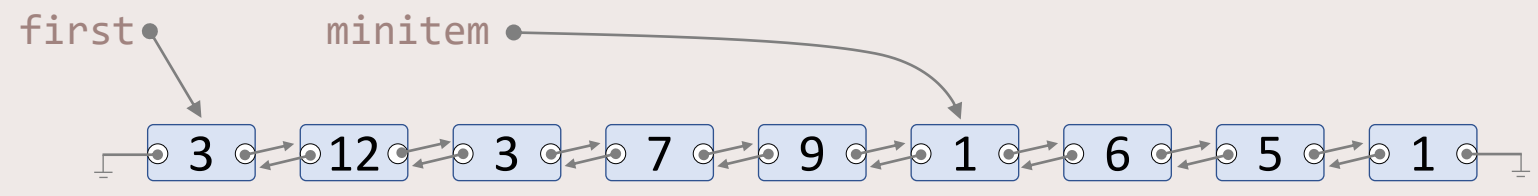


decreasekey is $O(1)$

decreasekey(*item*, *new key*)

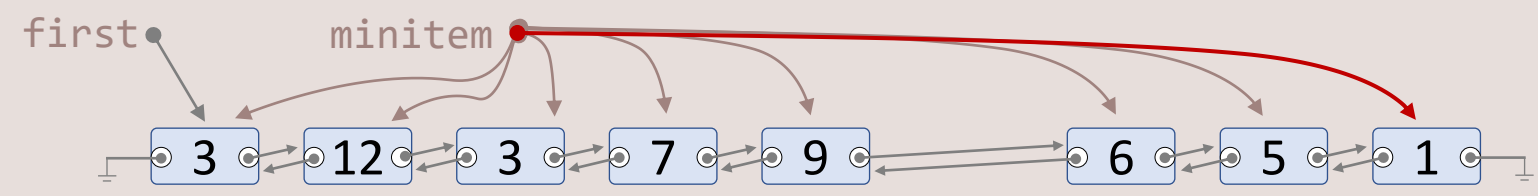
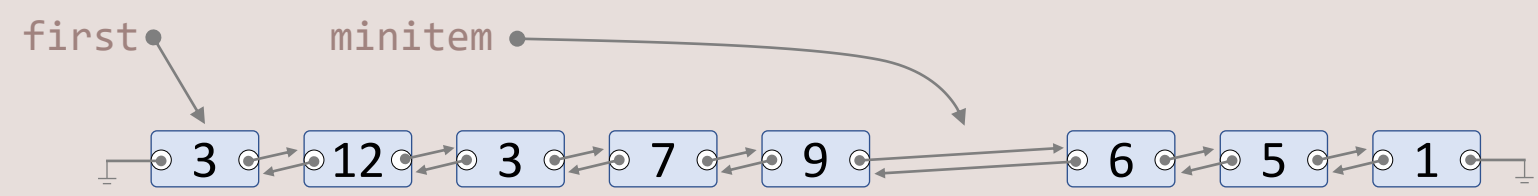


Linked-list priority queue



popmin is $O(N)$

popmin()



	popmin	push	decreasekey	
binary heap	$O(\log N)$	$O(\log N)$	$O(\log N)$	batch-push is $O(N)$
binomial heap	$O(\log N)$	$O(1)$ amort	$O(\log N)$	
linked list	$O(N)$	$O(1)$	$O(1)$	
Fibonacci heap	$O(\log N)$ amort	$O(1)$ amort	$O(1)$ amort	

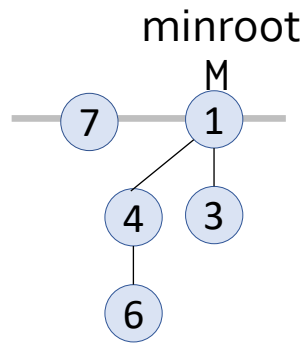
Design strategy for the Fibonacci heap:

- ❖ Give your data enough structure that you only need to touch a little bit of it
- ❖ Be lazy: let mess accumulate
- ❖ Do cleanup in batches



SECTION 7.6

The Fibonacci Heap



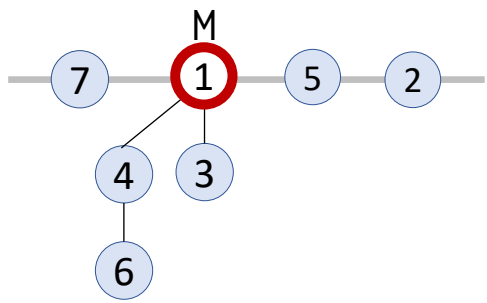
- store a list of trees, each a heap
- trees can have any shape
- keep track of the minroot

```

1 # Maintain a list of heaps (i.e. store a pointer to the root of each heap)
2 roots = []
3
4 # Maintain a pointer to the smallest root
5 minroot = None
6
7 def push(Value v, Key k):
8     create a new heap h consisting of a single item (v,k)
9     add h to the list of roots
10    update minroot if minroot is None or k < minroot.key

```

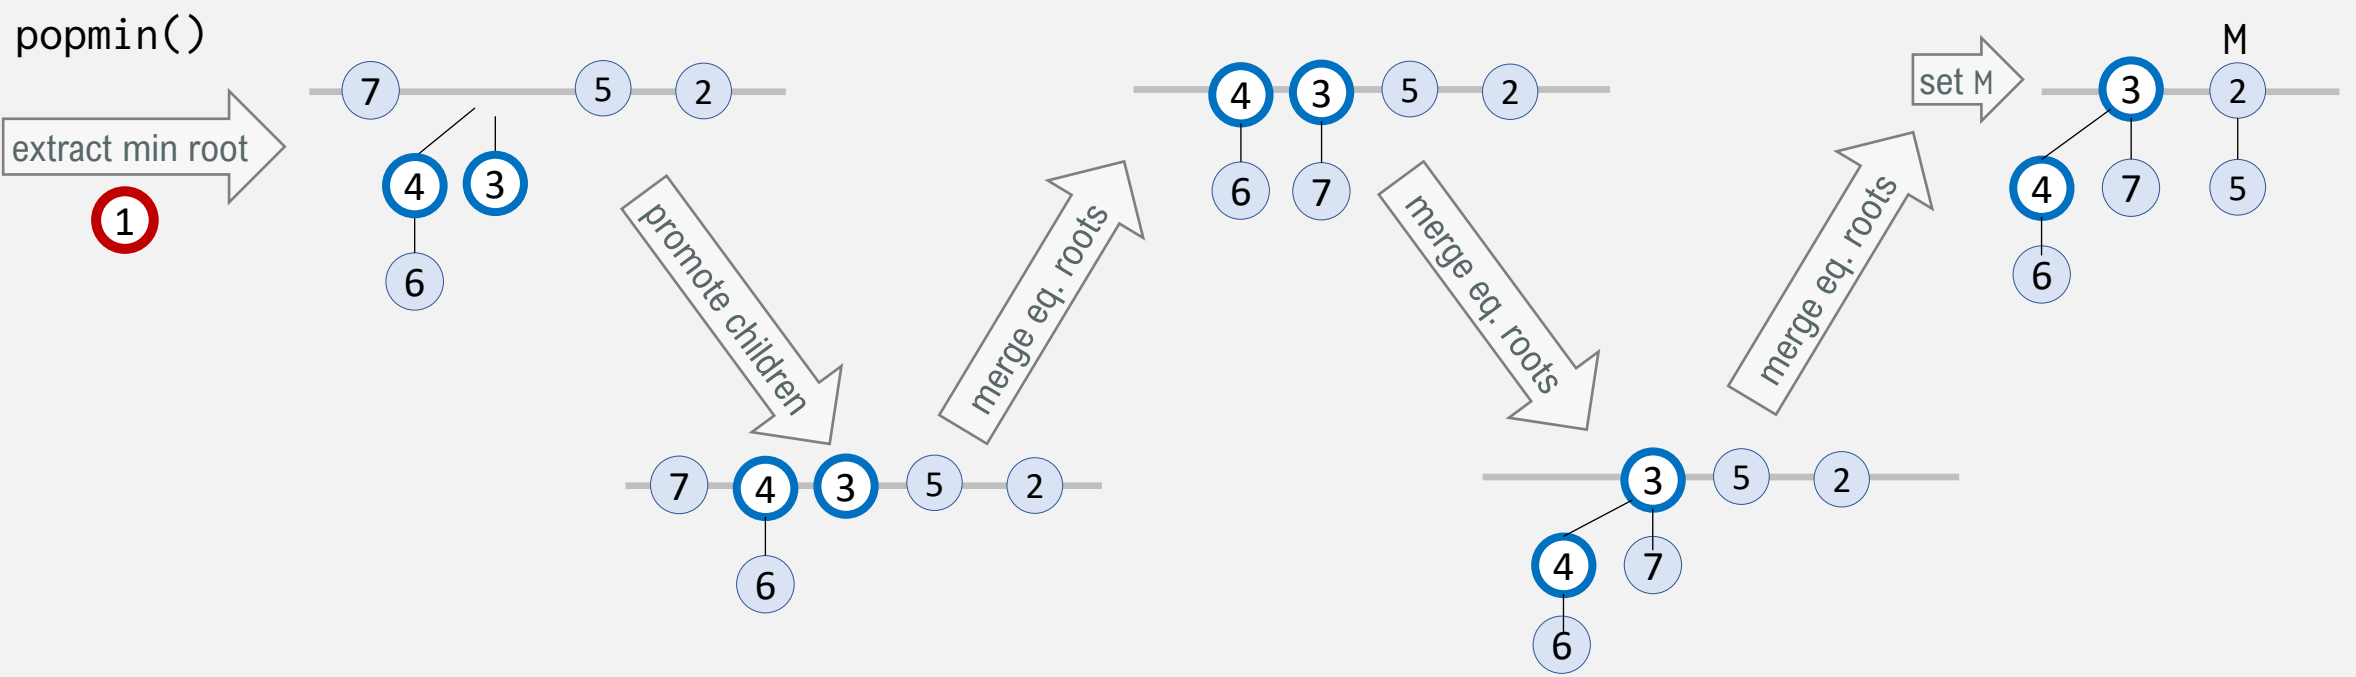




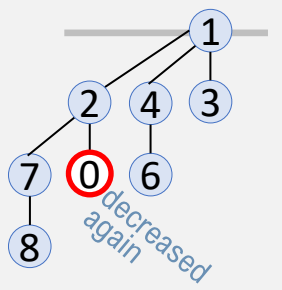
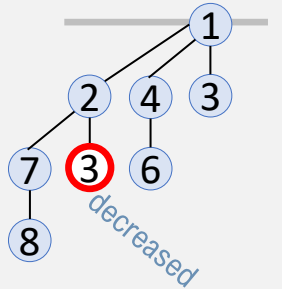
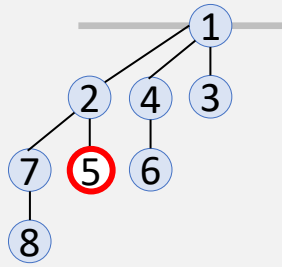
```

12 def popmin():
13     take note of minroot.value and minroot.key
14     delete the minroot node, and promote its children to be roots
15     # cleanup the roots
16     while there are two roots with the same degree:
17         merge those two roots, by making the larger root a child of the smaller
18     update minroot to point to the root with the smallest key
19     return the value and key we noted in line 13

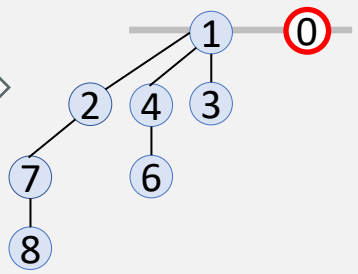
```



decreasekey(item, new key)



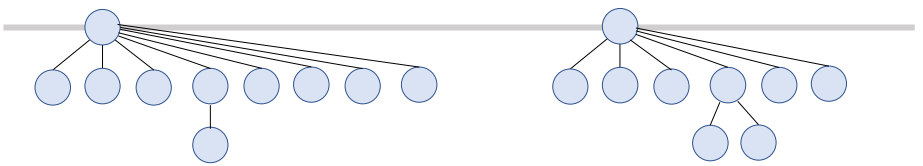
restore heap



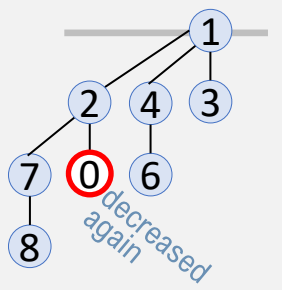
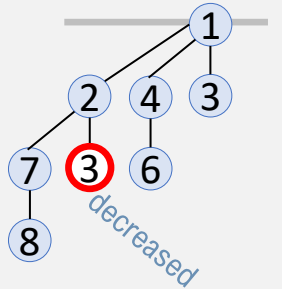
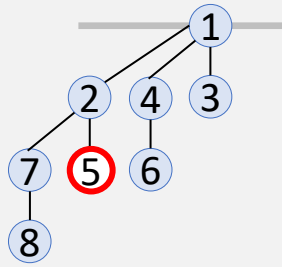
LAZY STRATEGY

Dump heap-violating nodes into the root list, to be cleaned up by the next popmin()

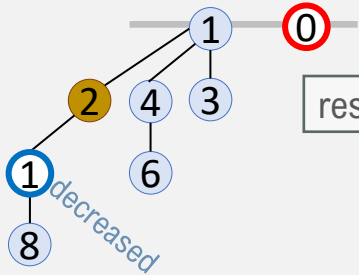
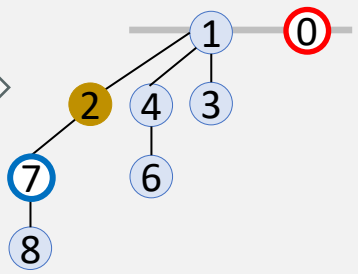
... **but** we might end up with a heap with wide shallow trees, which will make popmin() slow



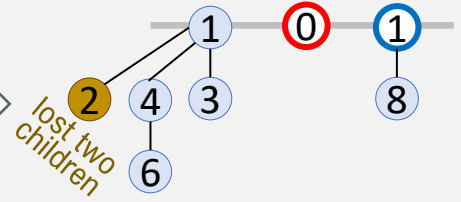
decreasekey(item, new key)



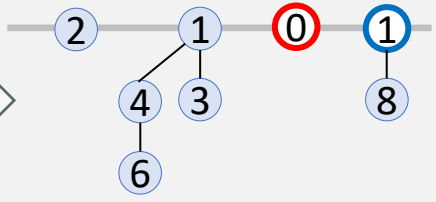
restore heap



restore heap



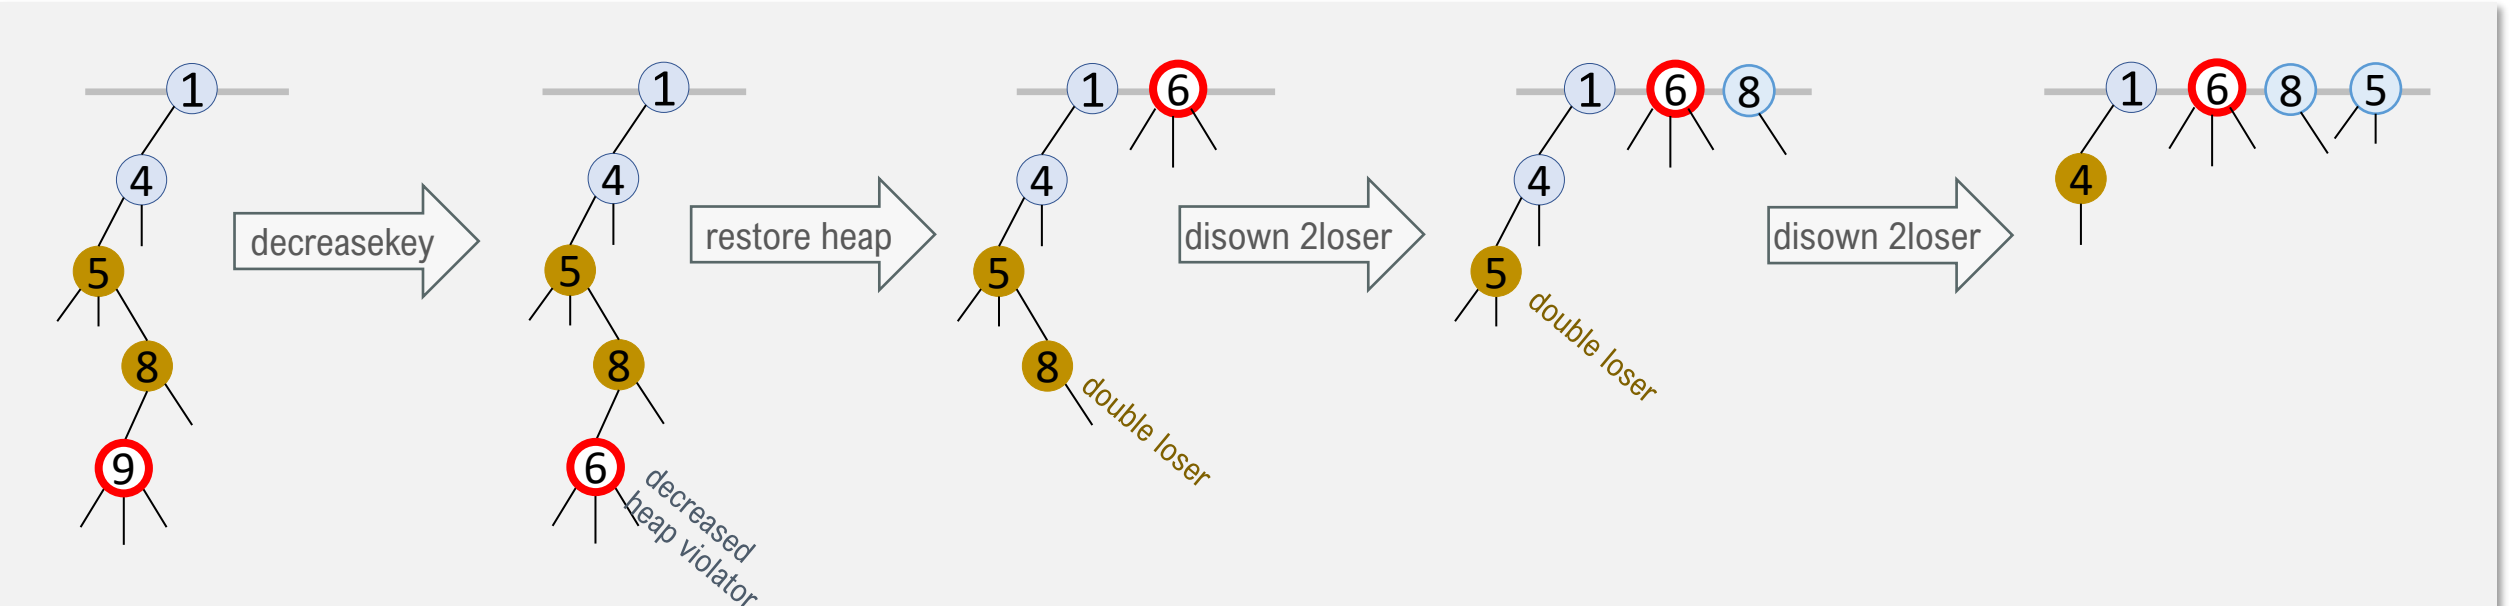
disown & unmark



Rule 1. Lose one child, and you're marked a **LOSER**

Rule 2. Lose two children, and you're dumped into the root list

```
30 # Every node will store a flag, n.loser = True / False
31
32 def decreasekey(v, k'):
33     let n be the node where this value is stored
34     n.key = k'
35     if n violates the heap condition:
36         repeat:
37             p = n.parent
38             remove n from p.children
39             insert n into the list of roots, updating minroot if necessary
40             n.loser = False
41             n = p
42         until p.loser == False
43         if p is not a root:
44             p.loser = True
45
46 # Modify popmin so that when we promote minroot's children, we erase any loser flags
```





Sometimes it pays
to let mess build up

Your parents want
lots of grandchildren*

* and they'll disown you if
you don't have enough



SECTION 7.8

Amortized analysis of the Fibonacci Heap

Take-away: this is an elegant use of potential functions to account for *two separate* unbounded-cost operations.

FIBONACCI HEAP COMPLEXITY ANALYSIS

COMPLEXITY ANALYSIS

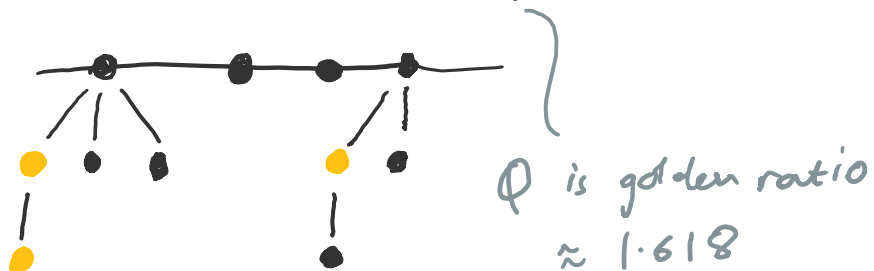
In a Fibonacci heap with N items,
using the potential function

$$\Phi = \text{num.roots} + 2 \times \text{num.losers},$$

- $\text{push}()$ has amortized cost $O(1)$
- $\text{decreasekey}()$ has amortized cost $O(1)$
- $\text{popmin}()$ has amortized cost $O(\log N)$

SHAPE THEOREM

Every node has degree $\leq \log_{\phi} N$



BINOMIAL HEAP COMPLEXITY ANALYSIS

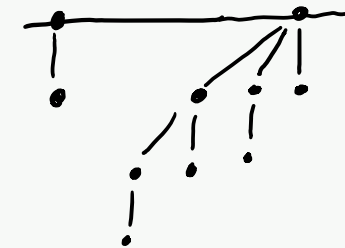
COMPLEXITY ANALYSIS

In a binomial heap with N items

- $\text{push}()$ is $O(\log N)$
- $\text{decreasekey}()$ is $O(\log N)$
- $\text{popmin}()$ is $O(\log N)$

SHAPE THEOREM

The largest tree has degree $\leq \log_2 N$



```
7 def push(Value v, Key k):  
8     create a new heap  $h$  consisting of a single item  $(v,k)$   
9     add  $h$  to the list of roots  
10    update minroot if minroot is None or  $k < \text{minroot.key}$ 
```

$$c = O(1)$$

$$\Delta \Phi = 1$$

am.cost

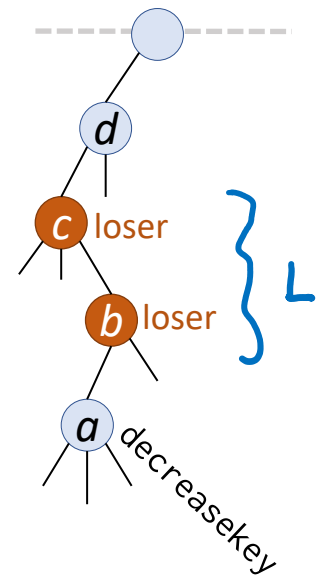
$$c' = c + \Delta \Phi = O(1)$$

$$\Phi = \text{num.roots} + 2 \times \text{num.losers}$$

```

32 def decreasekey(v, k'):
33     let n be the node where this value is stored
34     n.key = k'
35     if n violates the heap condition:
36         repeat:
37             p = n.parent
38             remove n from p.children
39             insert n into the list of roots, updating minroot if necessary
40             n.loser = False
41             n = p
42         until p.loser == False
43         if p is not a root:
44             p.loser = True

```



Case I: no heap violation. $c = O(1)$ $\Delta \Phi = 0$ $c + \Delta \Phi = O(1)$.

Case II: heap violation.

1. move a to root list. $c = O(1)$ $\Delta \Phi = 1$ or $\Delta \Phi = -1$ if a way loser. $c + \Delta \Phi = O(1)$

2. move up L losers $c = O(L)$ $\Delta \Phi = +L - 2L = -L$. $c + \Delta \Phi = O(1)$

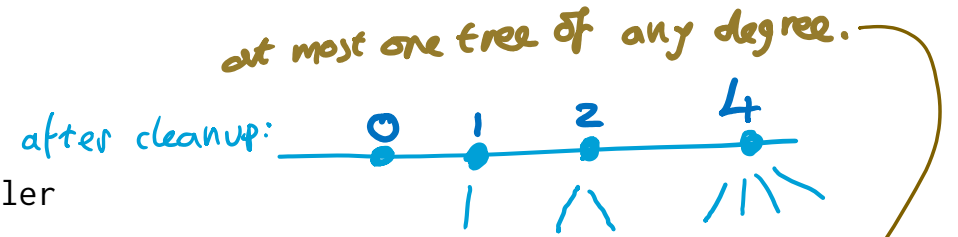
3. mark d as loser. $c = O(1)$ $\Delta \Phi = 2$ or $\Delta \Phi = 0$ if d was a root $c + \Delta \Phi = O(1)$

in each cases, $O(1)$ am. cost. } tot. is $O(1)$

```

12 def popmin():
13     take note of minroot.value and minroot.key
14     delete the minroot node, and promote its children to be roots
15     # cleanup the roots
16     while there are two roots with the same degree:
17         merge those two roots, by making the larger root a child of the smaller
18     update minroot to point to the root with the smallest key
19     return the value and key we noted in line 13

```



1. cut out minroot, promote its children

$$c = O(\# \text{children})$$

$$\frac{c + \Delta \Phi = O(\log N)}{\text{by Shape Theorem}}$$

$$\Delta \Phi \leq \# \text{children} - 1$$

* 2. cleanup. we'll see $c + \Delta \Phi = O(\log N)$

3. fix minroot.
by scanning all roots.

$$c = O(\log N)$$

$$\Delta \Phi = 0$$

$$\frac{c + \Delta \Phi = O(\log N)}{\text{by Shape Theorem}}$$

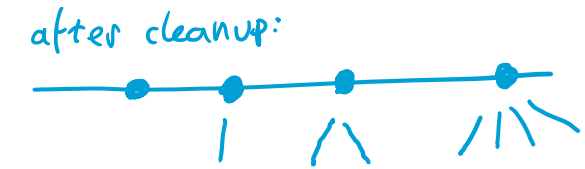
The total for these three steps is $O(\log N)$ amortized cost.

largest possible
root-degree
is $\lfloor \log_{\phi} N \rfloor$

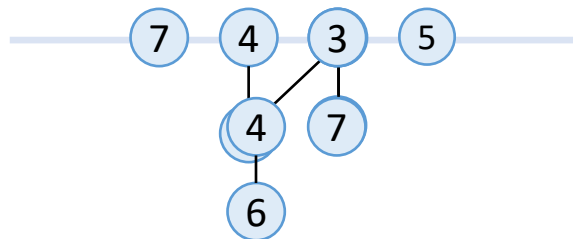
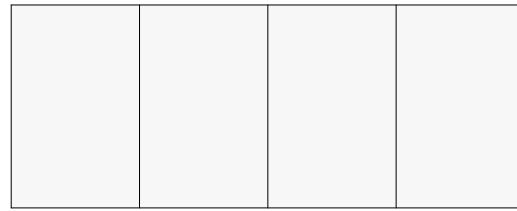
trees
after cleanup is
 $\leq 1 + \lfloor \log_{\phi} N \rfloor$



```
12 def popmin():
13     take note of minroot.value and minroot.key
14     delete the minroot node, and promote its children to be roots
15     # cleanup the roots
16     while there are two roots with the same degree:
17         merge those two roots, by making the larger root a child of the smaller
18     update minroot to point to the root with the smallest key
19     return the value and key we noted in line 13
```



```
20 def cleanup(roots):
21     root_array = [None, None, ....]
22     for each tree t in roots:
23         x = t
24         while root_array[x.degree] is not None:
25             u = root_array[x.degree]
26             root_array[x.degree] = None
27             x = merge(x, u)
28         root_array[x.degree] = x
29     roots = list of non-None values from root_array
```

for each t in roots:array of size $\lfloor \log_{\phi} N \rfloor + 1$.

root_array

updated roots:



Suppose I started with x trees, do M merges, end up with y trees, so $y = x - M \Leftrightarrow x = y + M$.

$$c = \underbrace{O(x)}_{\text{iteration over roots.}} + \underbrace{M}_{\text{merging}} + \underbrace{\log N}_{\text{copy over}} = \underbrace{O(y + 2M + \log N)}_{\text{since } x = y + M} = \underbrace{O(2M + 2 \log N)}_{\text{since } y \leq \lfloor \log_{\phi} N \rfloor + 1} = \underbrace{O(M + \log N)}_{\text{since constants don't matter for } O.}$$

$$\Delta \Phi = -M \quad \text{since } \# \text{ roots has decreased by } M$$

$$\text{Thus } c + \Delta \Phi = O(M + \log N) - M = O(\log N).$$

def cleanup(roots):

root_array = [None, None, ...]

for each tree t in roots: $x = t$

while root_array[x.degree] is not None:

 $u = \text{root_array}[x.\text{degree}]$

root_array[x.degree] = None

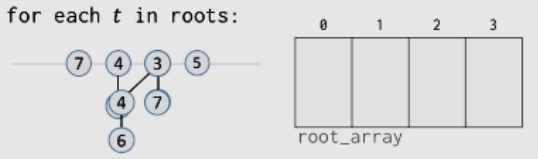
 $x = \text{merge}(x, u)$ root_array[x.degree] = u

roots = list of non-None values from root_array

After cleanup, we have at most one tree of each degree; max degree $\leq \lfloor \log_{\phi} N \rfloor$, so we may have up to $1 + \lfloor \log_{\phi} N \rfloor$ trees.

$\Phi = \text{num.roots} + 2 \times \text{num.losers}$ *pays in advance for these "uncontrolled" iterations*

for each t in roots:



updated roots:

Suppose we start with x trees, do M merges, and end up with y trees.

$$c = O(x + M + \log N) = O(y + 2M + \log N) = O(2M + 2\log N) = O(M + \log N)$$

$\Delta\Phi = -M$

am. cost is $c + \Delta\Phi = O(M + \log N) - M = O(\log N)$


```

20 def cleanup(roots):
21     root_array = [None, None, ...] ← empty array of size ⌊log N⌋ + 1
22     for each tree t in roots:
23         x = t
24         while root_array[x.degree] is not None:
25             u = root_array[x.degree]
26             root_array[x.degree] = None
27             x = merge(x, u)
28             root_array[x.degree] = u
29     roots = list of non-None values from root_array
    
```

popmin
had to do M merges

```

32 def decreasekey(v, k'):
33     let n be the node where this value is stored
34     n.key = k'
35     if n violates the heap condition:
36         repeat:
37             p = n.parent
38             remove n from p.children
39             insert n into the list of roots, updating minroot if necessary
40             n.loser = False
41             n = p
42         until p.loser == False
43         if p is not a root:
44             p.loser = True
    
```



CASE I: no heap violation
 $c = O(1)$ $\Delta\Phi = 0 \Rightarrow c + \Delta\Phi = O(1)$

CASE II: heap violation

1. move a to rootlist
 $c = O(1)$ $\Delta\Phi = 1$ or $\Delta\Phi = -1$ if a was loser $\Rightarrow c + \Delta\Phi = O(1)$
2. Move up L losers also
 $c = O(L)$ $\Delta\Phi = +L - 2L = -L \Rightarrow c + \Delta\Phi = O(1)$
3. mark d as a loser unless d is root, $\Delta\Phi = 0$
 $c = O(1)$ $\Delta\Phi = 2 \Rightarrow c + \Delta\Phi = O(1)$

in both cases, total amortized cost is $O(1)$

decreasekey
had to move L nodes to root