



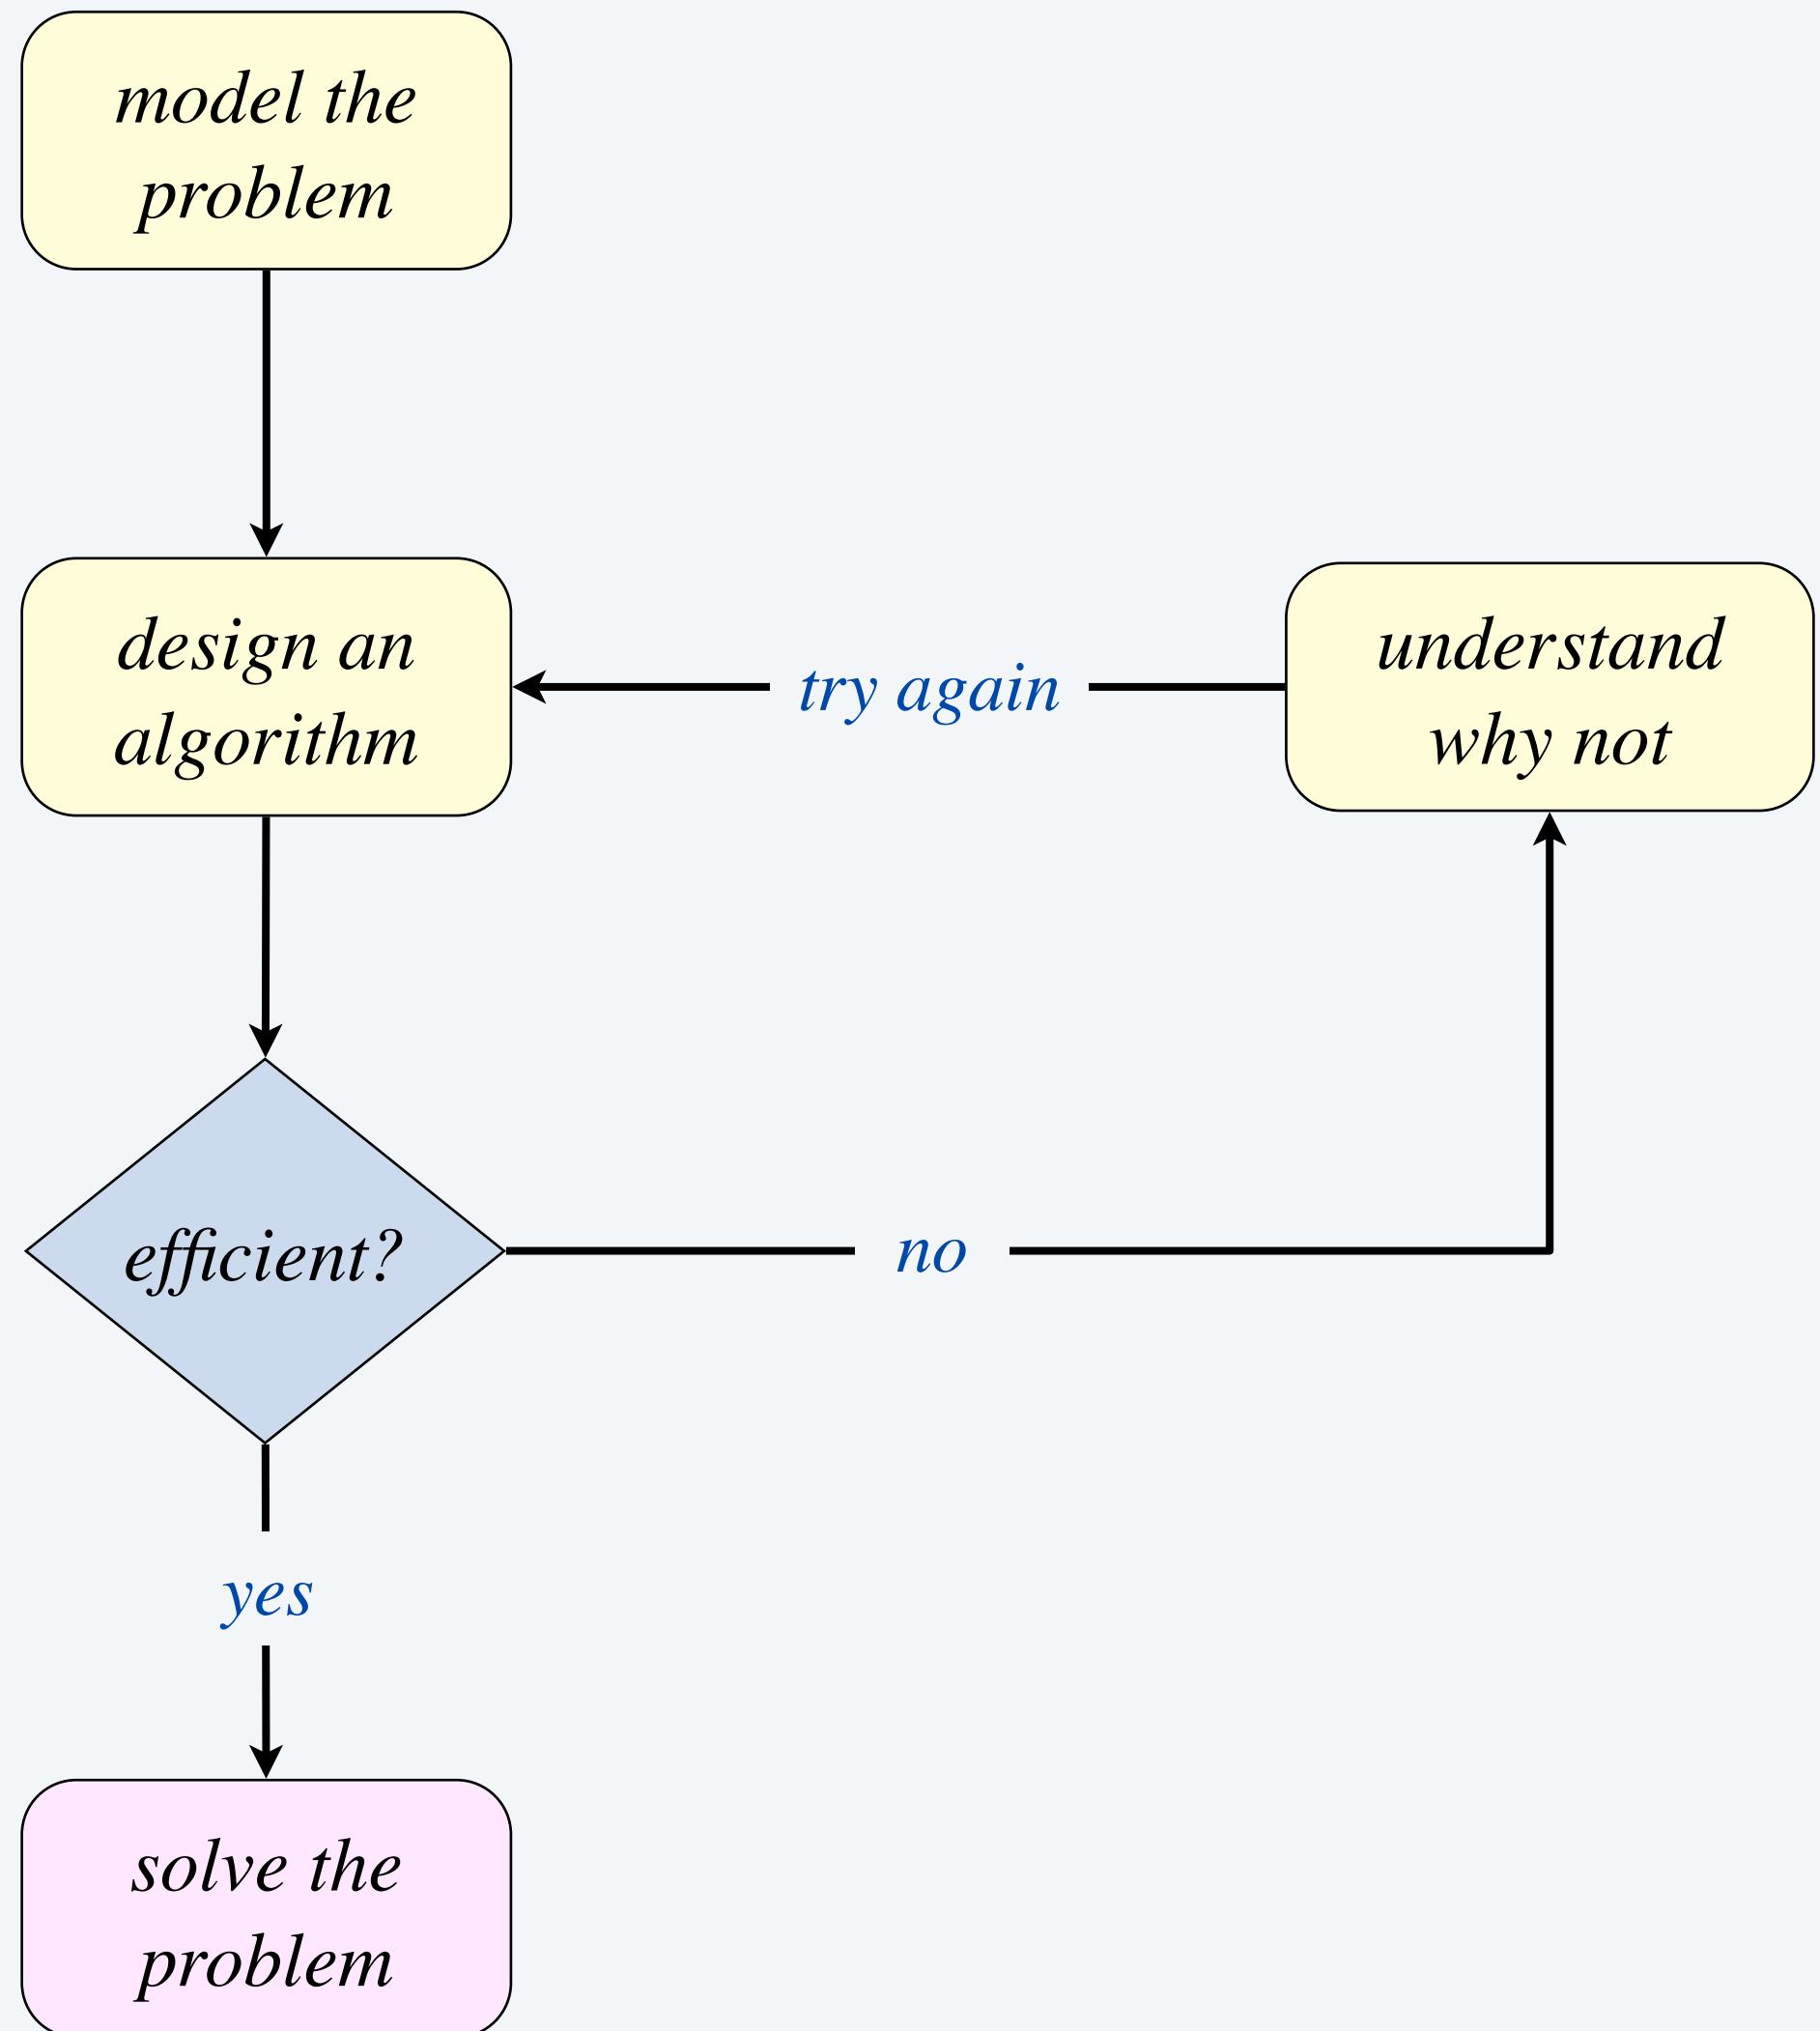
1.5 UNION-FIND

- ▶ *union-find data type*
- ▶ *quick-find*
- ▶ *quick-union*
- ▶ *weighted quick-union*

<https://algs4.cs.princeton.edu>

Subtext of today's lecture (and this course)

Steps to develop a usable algorithm to solve a computational problem.



1.5 UNION-FIND

- ▶ *union-find data type*
- ▶ *quick-find*
- ▶ *quick-union*
- ▶ *weighted quick-union*
- ▶ *percolation*



Algorithms

ROBERT SEDGEWICK | KEVIN WAYNE

<https://algs4.cs.princeton.edu>

Union-find data type

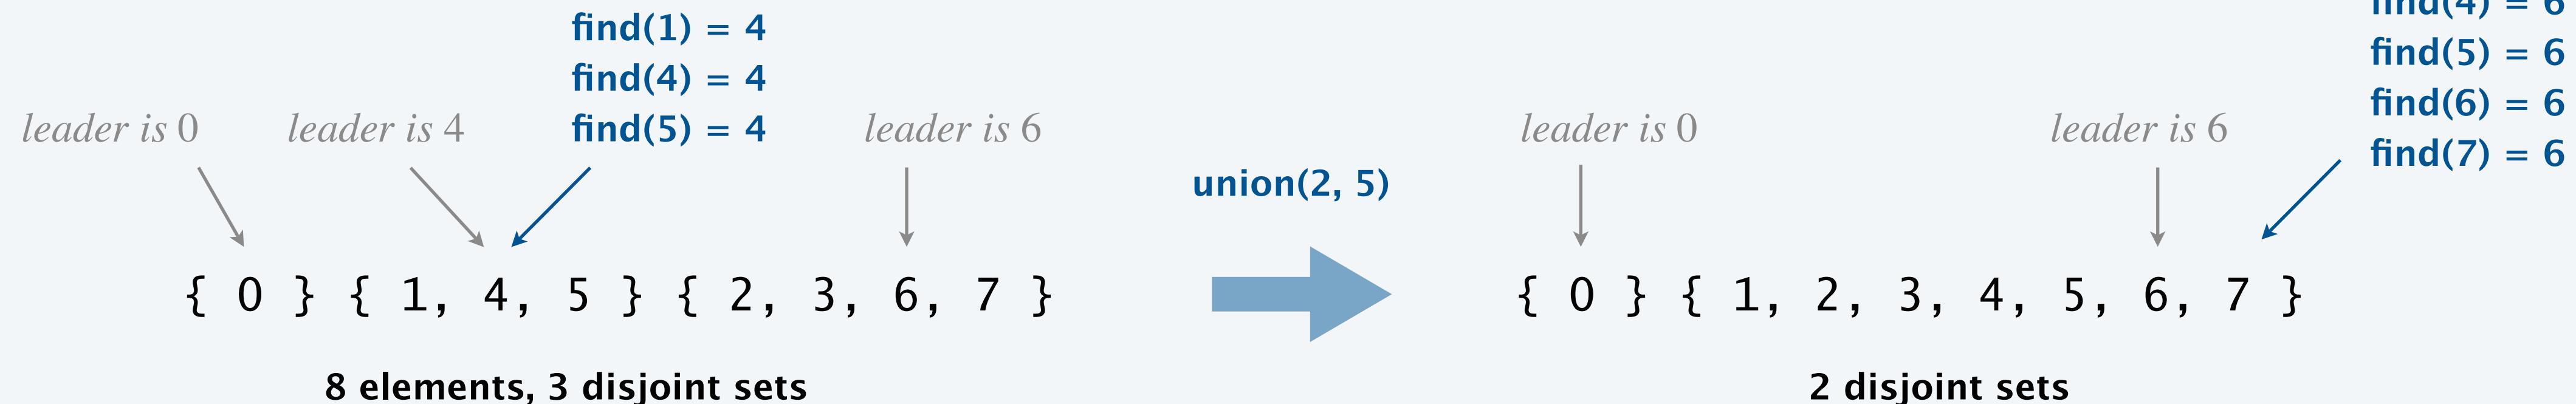
Disjoint sets. A collection of sets containing n elements, with each element in exactly one set.

Leader. Each set designates one of its elements as **leader** (to uniquely identify it).

*no restriction on which element is designated leader
(but leader of a set can't change unless the set changes)*

Find. Return the leader of the set containing element p . ← *main use case:
are two elements in the same set ?*

Union. Merge the set containing element p with the set containing element q .



Union-find data type: API

Goal. Design an **efficient** union-find data type.

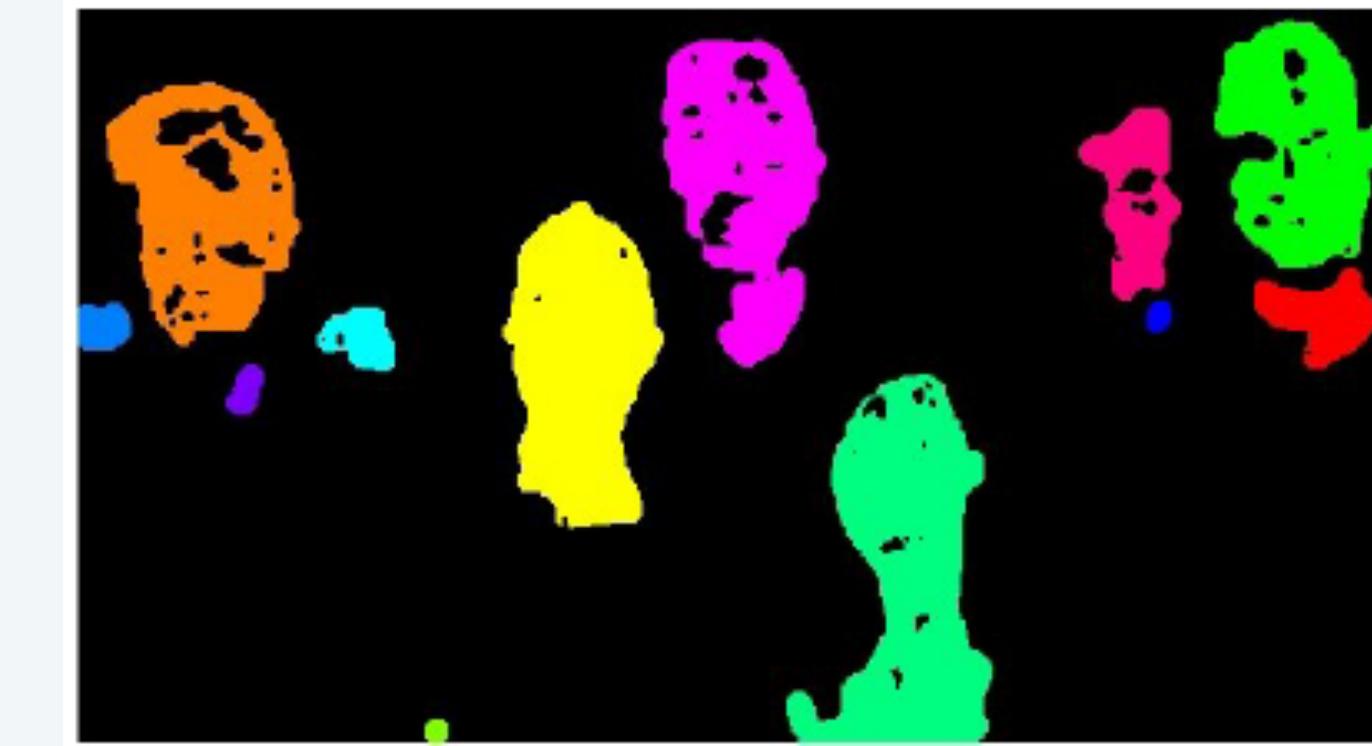
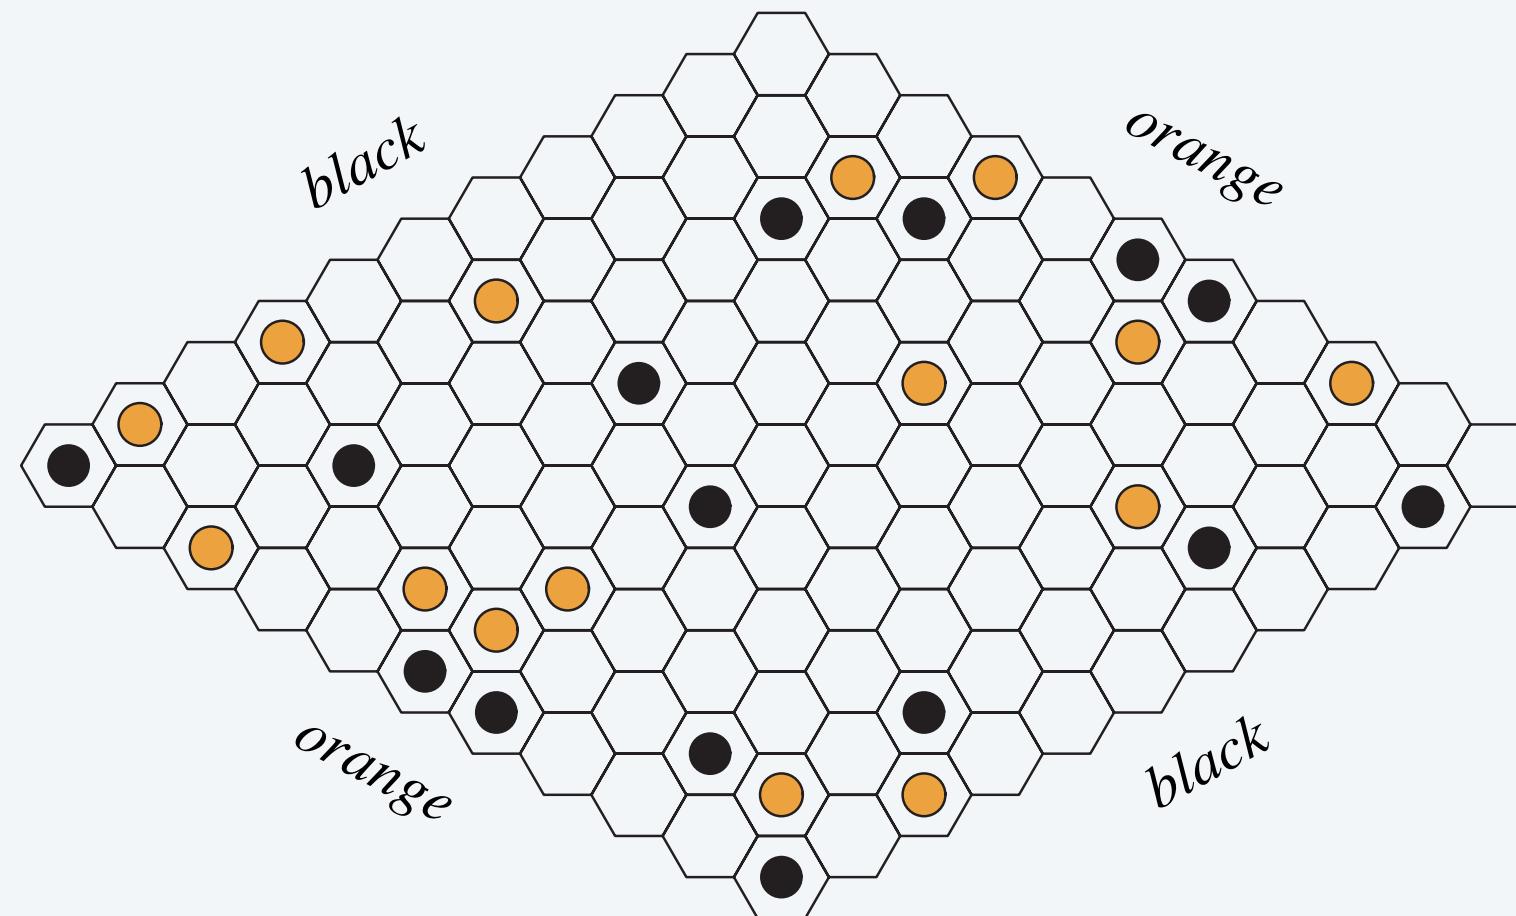
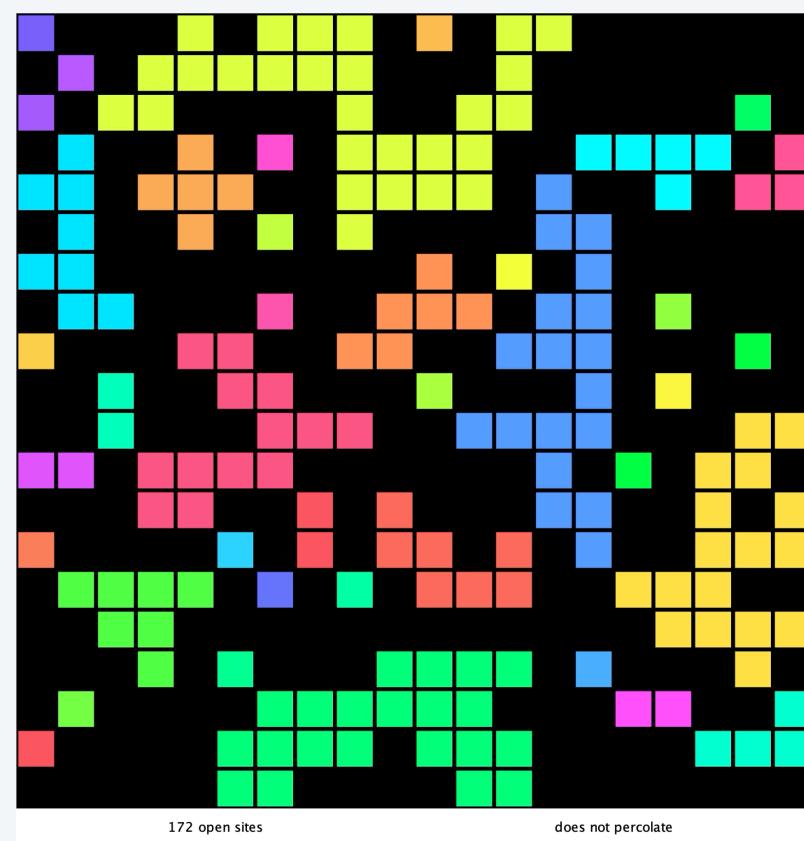
- Simplifying assumption: the n elements are named $0, 1, 2, \dots, n - 1$.
- The `union()` and `find()` operations can be intermixed.
- Number of elements n can be huge.
- Number of operations m can be huge.

public class UF	description
<code>UF(int n)</code>	<i>initialize with n singleton sets (0 to $n - 1$)</i>
<code>void union(int p, int q)</code>	<i>merge sets containing elements p and q</i>
<code>int find(int p)</code>	<i>return the leader of set containing element p</i>

Union-find data type: applications

Disjoint sets can represent:

- Clusters of conducting sites in a composite system. ← *see Assignment 1 (Percolation)*
- Connected components in a graph. ← *see Kruskal's algorithm (MST lecture)*
- Interlinked friends in a social network.
- Interconnected devices in a mobile network.
- Equivalent variable names in a Fortran program.
- Adjoining stones of the same color in the game of Hex.
- Contiguous pixels corresponding to same feature in a digital image.





1.5 UNION-FIND

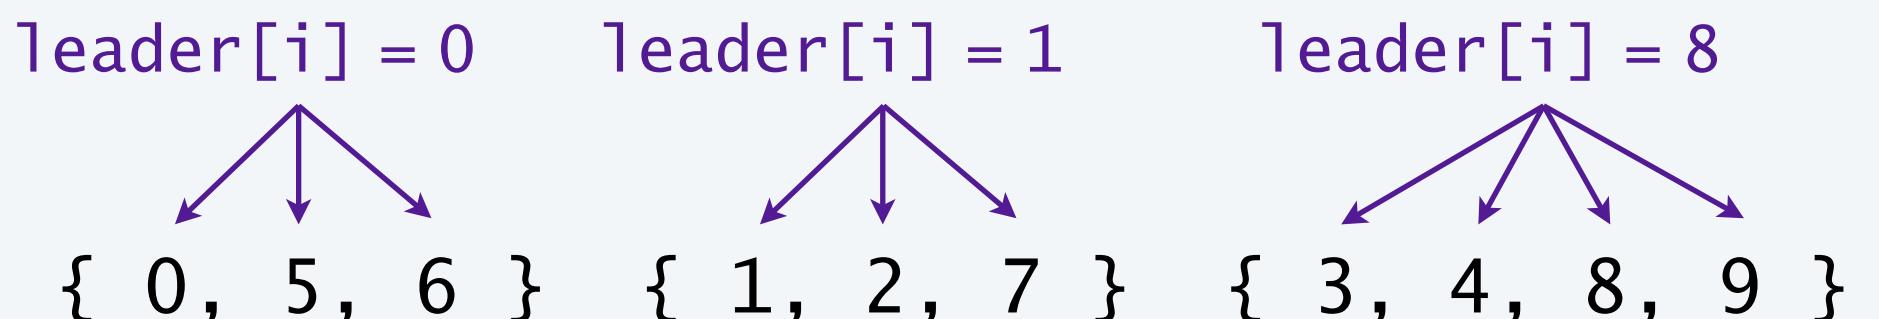
- ▶ *union-find data type*
- ▶ **quick-find**
- ▶ **quick-union**
- ▶ *weighted quick-union*

Quick-find

Data structure.

- Integer array `leader[]` of length n .
- Interpretation: `leader[i]` is the leader of the set containing element i .

	0	1	2	3	4	5	6	7	8	9
<code>leader[]</code>	0	1	1	8	8	0	0	1	8	8



10 elements, 3 disjoint sets

Q. How to implement `find(p)`?

A. Easy, just return `leader[p]`.

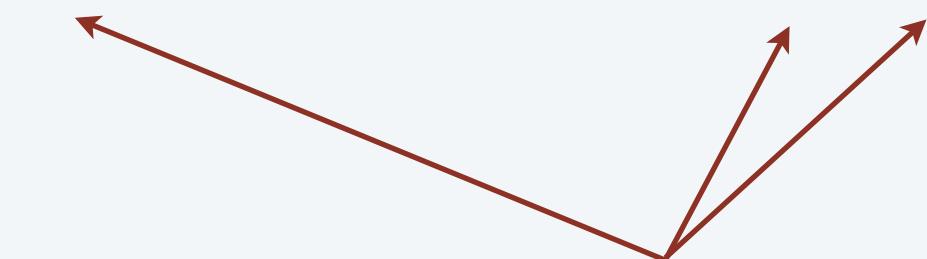
Quick-find

Data structure.

- Integer array `leader[]` of length n .
- Interpretation: `leader[i]` is the leader of the set containing element i .

`union(6, 2)`

	0	1	2	3	4	5	6	7	8	9
<code>leader[]</code>	1	1	1	8	8	1	1	1	8	8



*performance issue:
many array elements can change*

Q. How to implement `union(p, q)`?

A. Change all array elements whose value is `leader[p]` to `leader[q]`. \longleftrightarrow *or vice versa*

Quick-find: Java implementation

```
public class QuickFindUF {  
    private int[] leader;  
  
    public QuickFindUF(int n) {  
        leader = new int[n];  
        for (int i = 0; i < n; i++)  
            leader[i] = i;  
    }  
  
    public int find(int p) {  
        return leader[p];  
    }  
  
    public void union(int p, int q) {  
        int leaderP = leader[p];  
        int leaderQ = leader[q];  
        for (int i = 0; i < leader.length; i++) ←  
            if (leader[i] == leaderP)  
                leader[i] = leaderQ;  
    }  
}
```

← *initialize leader of each element to itself
(n array accesses)*

← *return the leader of p
(1 array access)*

← *change all array elements whose
value is $\text{leader}[p]$ to $\text{leader}[q]$
($\geq n$ array accesses)*

Quick-find is too slow

Cost model. Number of array accesses (for read or write).

algorithm	initialize	union	find
quick-find	n	n	1

worst-case number of array accesses (ignoring leading coefficient)

Union is too expensive. Processing any sequence of m `union()` operations on n elements takes $\geq mn$ array accesses.


quadratic in input size!

Ex. Performing 10^9 `union()` operations on 10^9 elements might take 30 years.



1.5 UNION-FIND

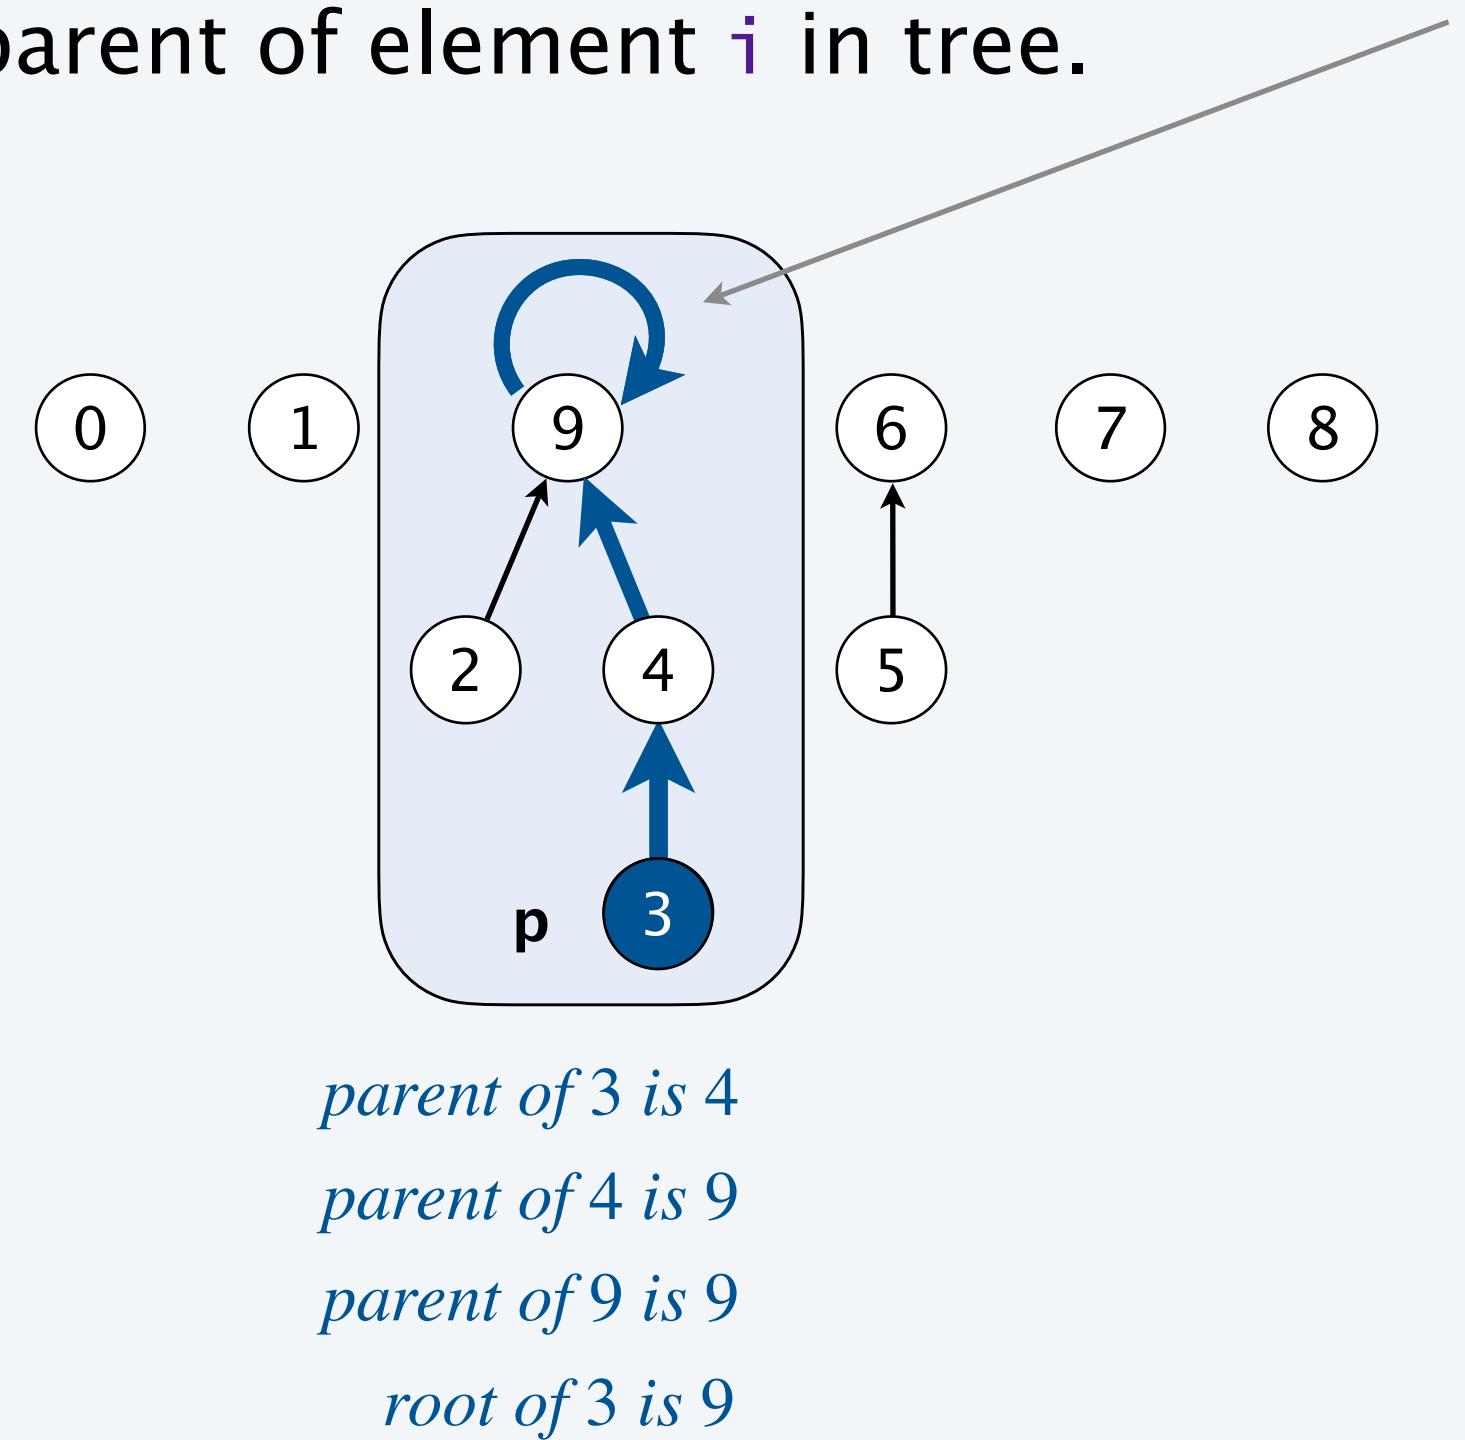
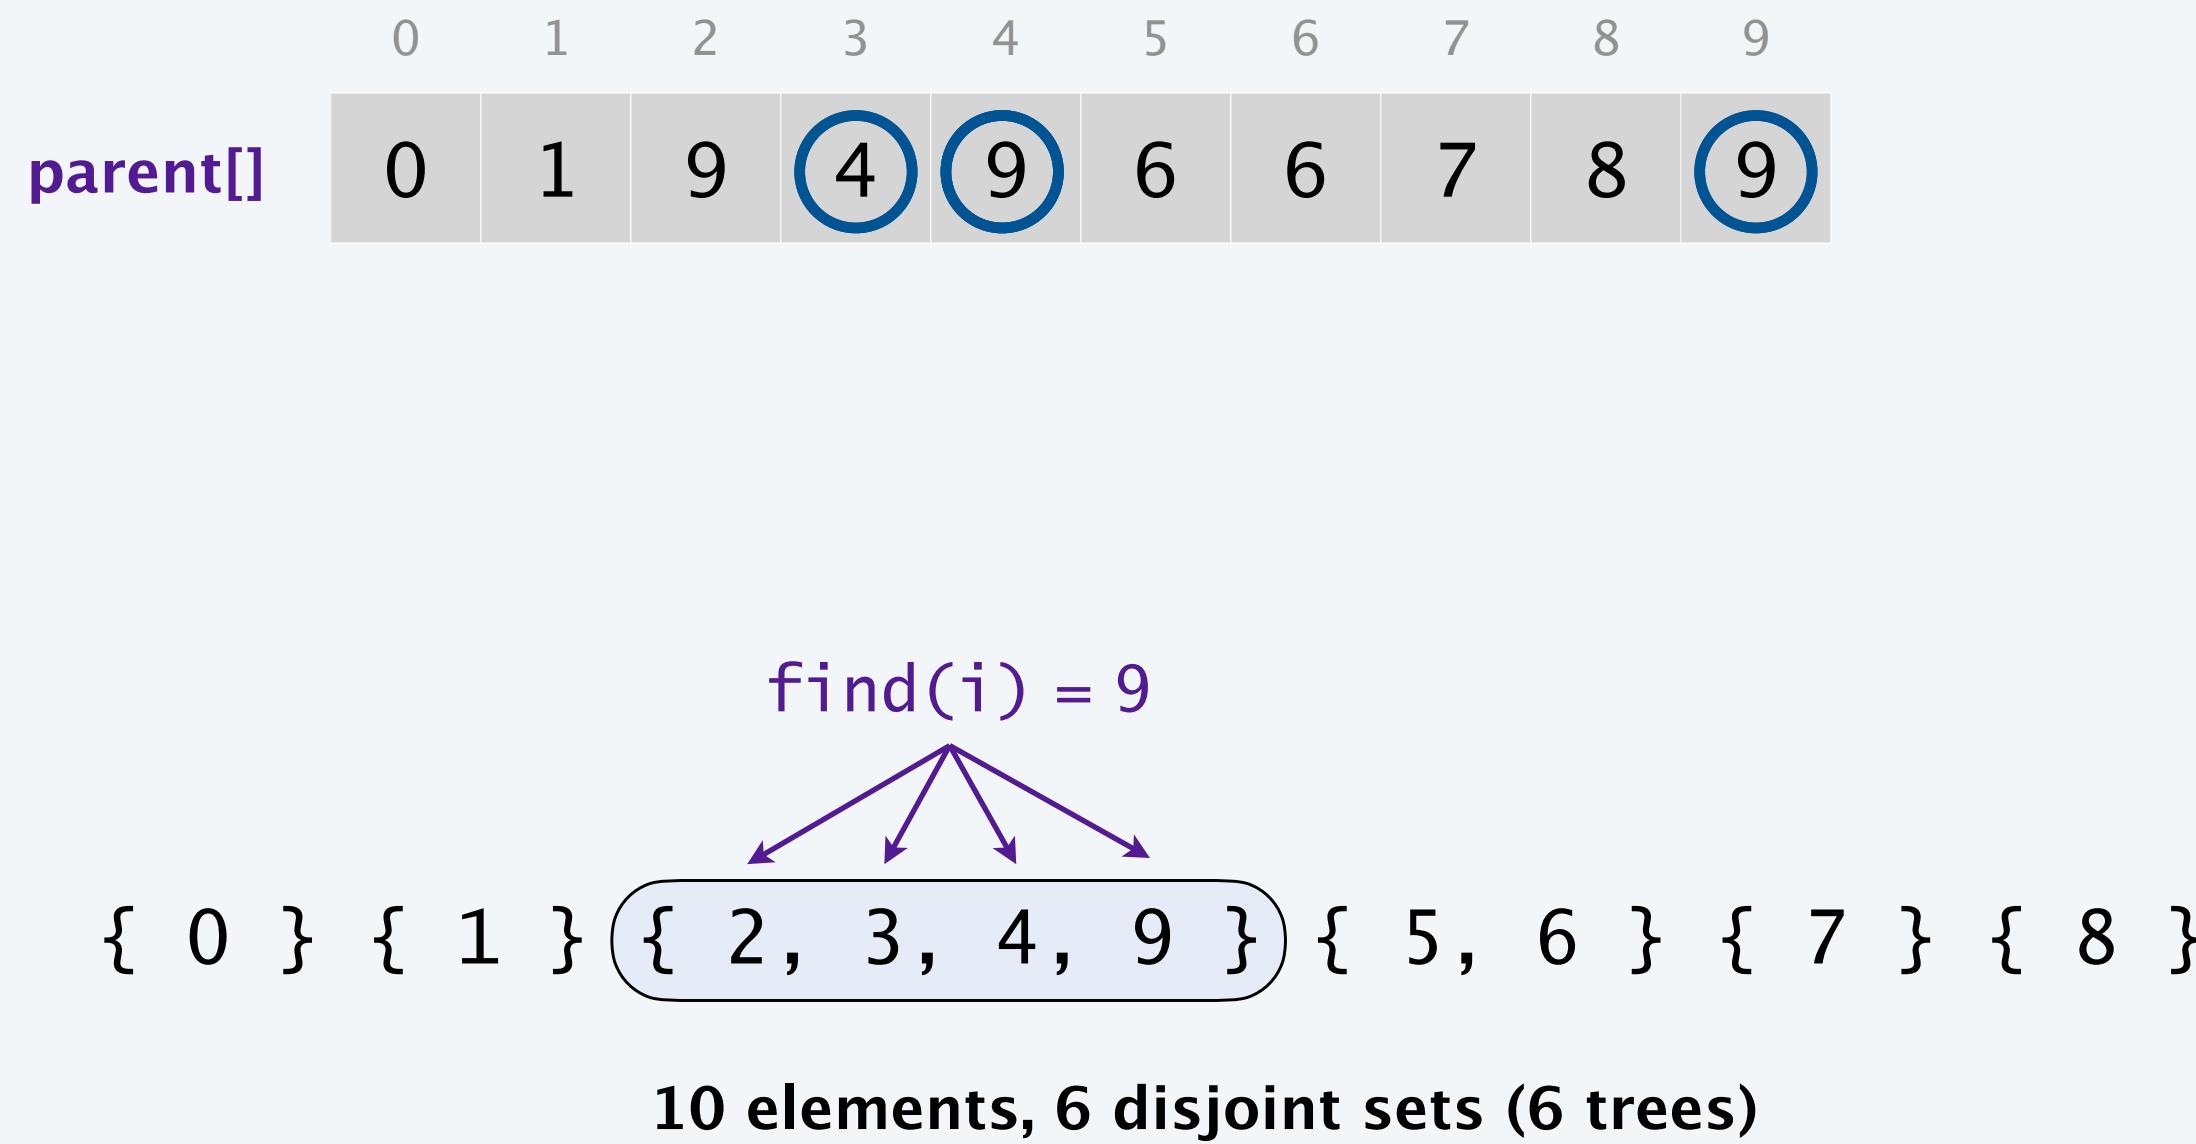
- ▶ *union-find data type*
- ▶ *quick-find*
- ▶ *quick-union*
- ▶ *weighted quick-union*

Quick-union

Data structure: Forest-of-trees.

- Interpretation: elements in one rooted tree correspond to one set.
- Integer array `parent[]` of length `n`, where `parent[i]` is parent of element `i` in tree.

*by convention, we make root point to itself
(but typically suppress from drawing)*



Q. How to implement `find(p)`?

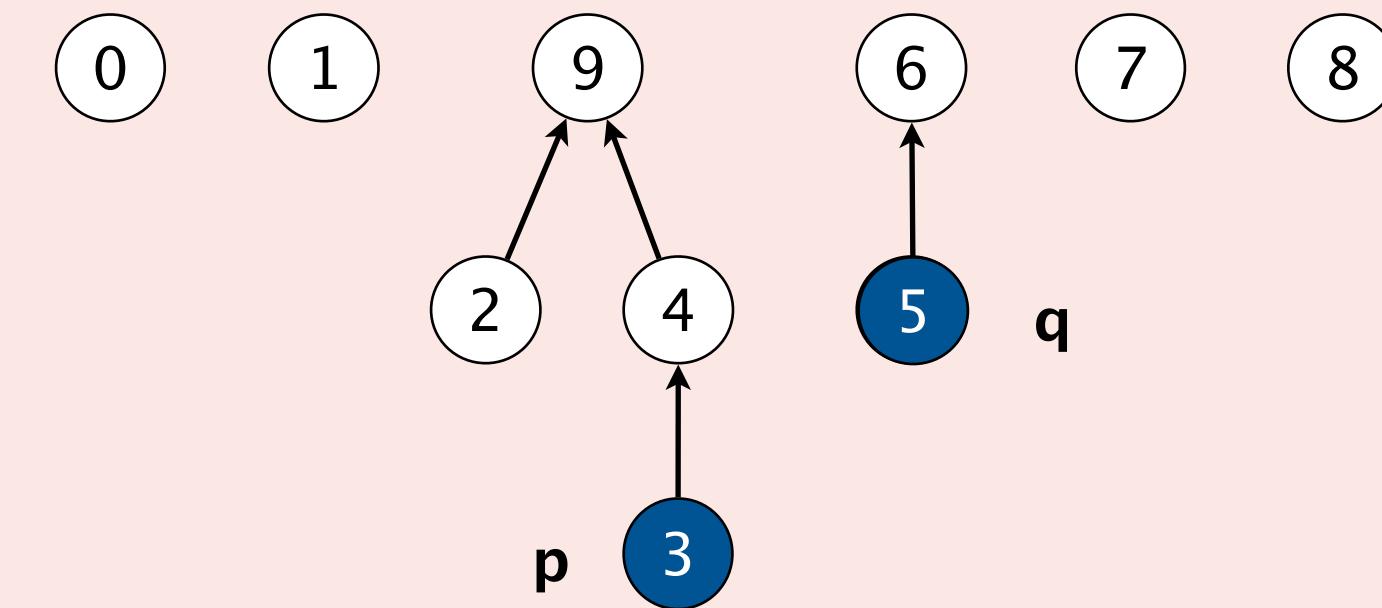
A. Use tree roots as leaders \Rightarrow return root of tree containing `p`.



Data structure: Forest-of-trees.

- Interpretation: elements in one rooted tree correspond to one set.
- Integer array `parent[]` of length `n`, where `parent[i]` is parent of element `i` in tree.

0	1	2	3	4	5	6	7	8	9	
<code>parent[]</code>	0	1	9	4	9	6	6	7	8	9



Which is **not** a valid way to implement `union(3, 5)` ?

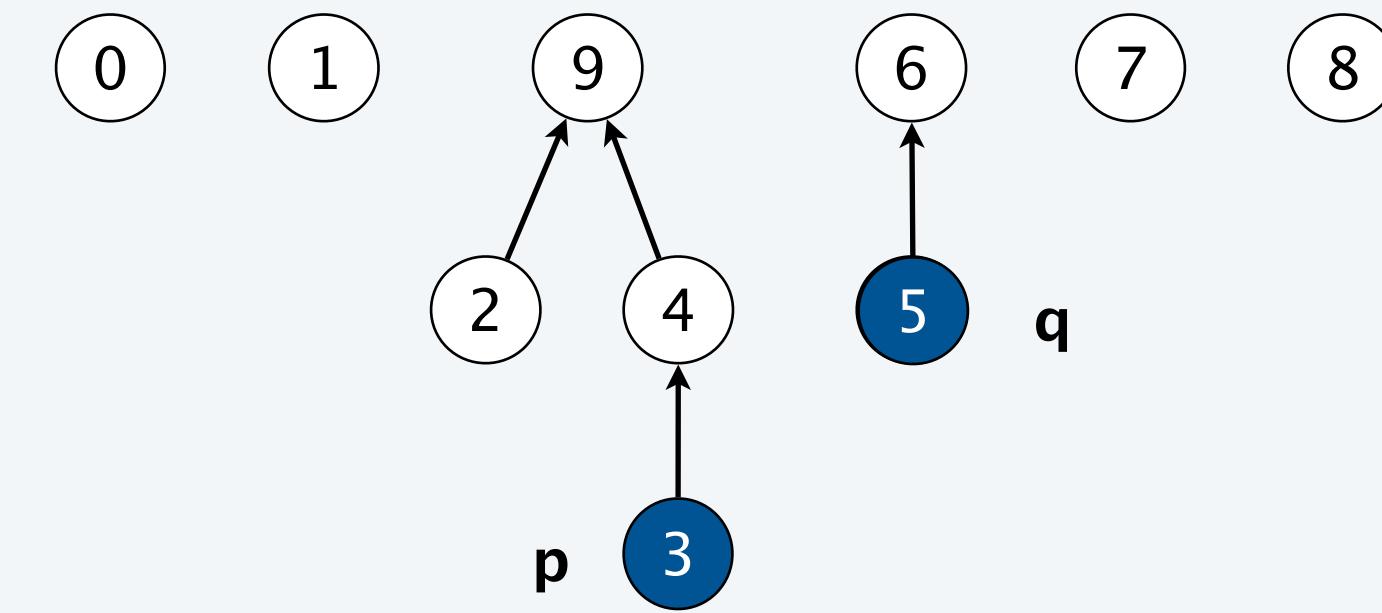
- A. Set `parent[6] = 9`.
- B. Set `parent[9] = 6`.
- C. Set `parent[3] = 5`.
- D. Set `parent[2] = parent[3] = parent[4] = parent[9] = 6`.

Quick-union

Data structure: Forest-of-trees.

- Interpretation: elements in one rooted tree correspond to one set.
- Integer array `parent[]` of length `n`, where `parent[i]` is parent of element `i` in tree.

	0	1	2	3	4	5	6	7	8	9
<code>union(3, 5)</code>	0	1	9	4	9	6	6	7	8	9



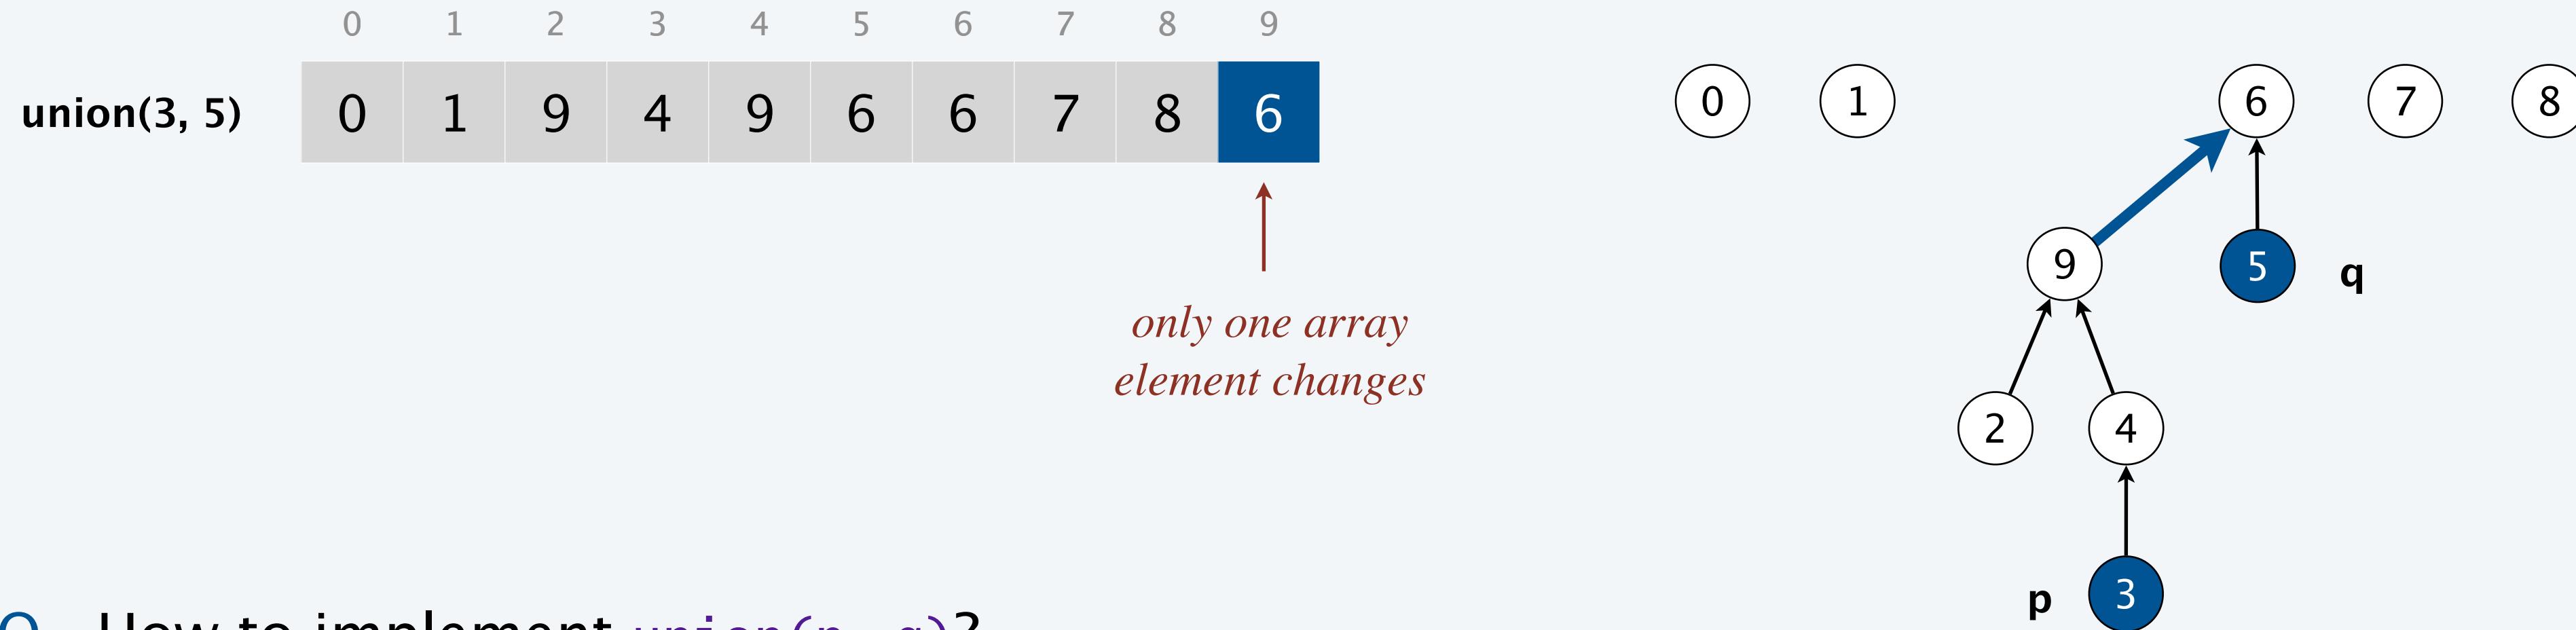
Q. How to implement `union(p, q)`?

A. Set `parent[p's root] = q's root.` \leftarrow or vice versa

Quick-union

Data structure: Forest-of-trees.

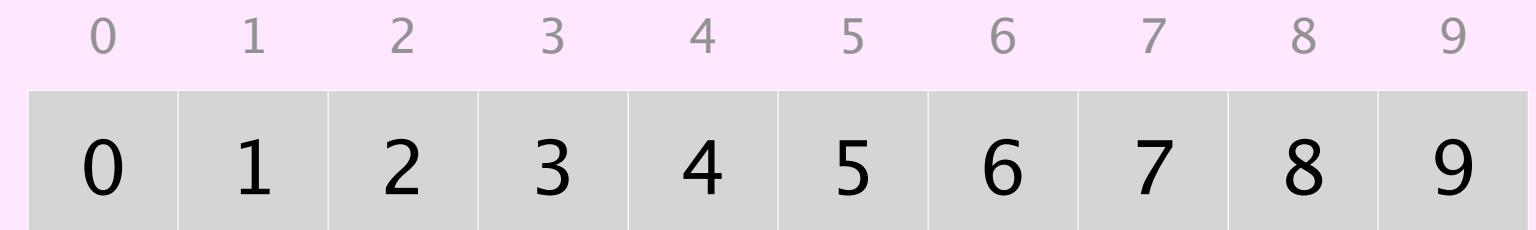
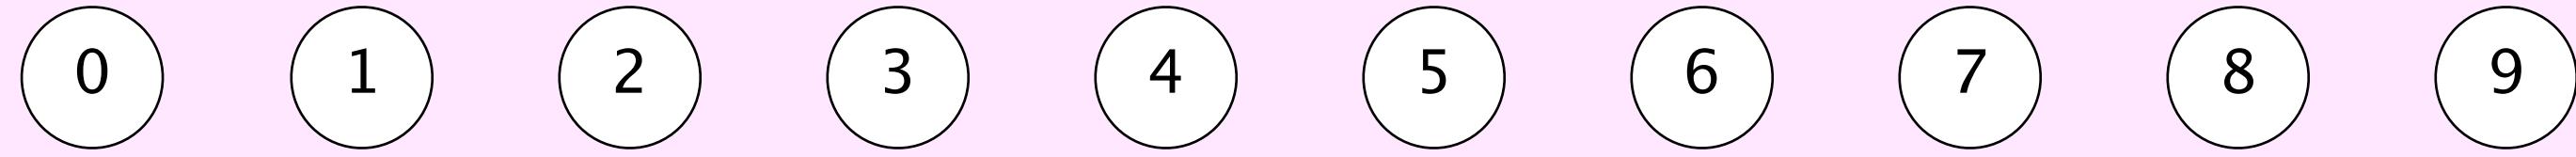
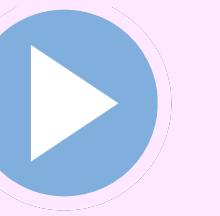
- Interpretation: elements in one rooted tree correspond to one set.
- Integer array `parent[]` of length `n`, where `parent[i]` is parent of element `i` in tree.



Q. How to implement `union(p, q)`?

A. Set `parent[p's root] = q's root.` \leftarrow or vice versa

Quick-union demo



Quick-union: Java implementation

```
public class QuickUnionUF {  
    private int[] parent;  
  
    public QuickUnionUF(int n) {  
        parent = new int[n];  
        for (int i = 0; i < n; i++)  
            parent[i] = i;  
    }  
  
    public int find(int p) {  
        while (p != parent[p])  
            p = parent[p];  
        return p;  
    }  
  
    public void union(int p, int q) {  
        int rootP = find(p);  
        int rootQ = find(q);  
        parent[rootP] = rootQ;  
    }  
}
```

*set parent of each element to itself
(to create forest of n singleton trees)*

*follow parent pointers until reach root;
return resulting root*

link root of p to root of q

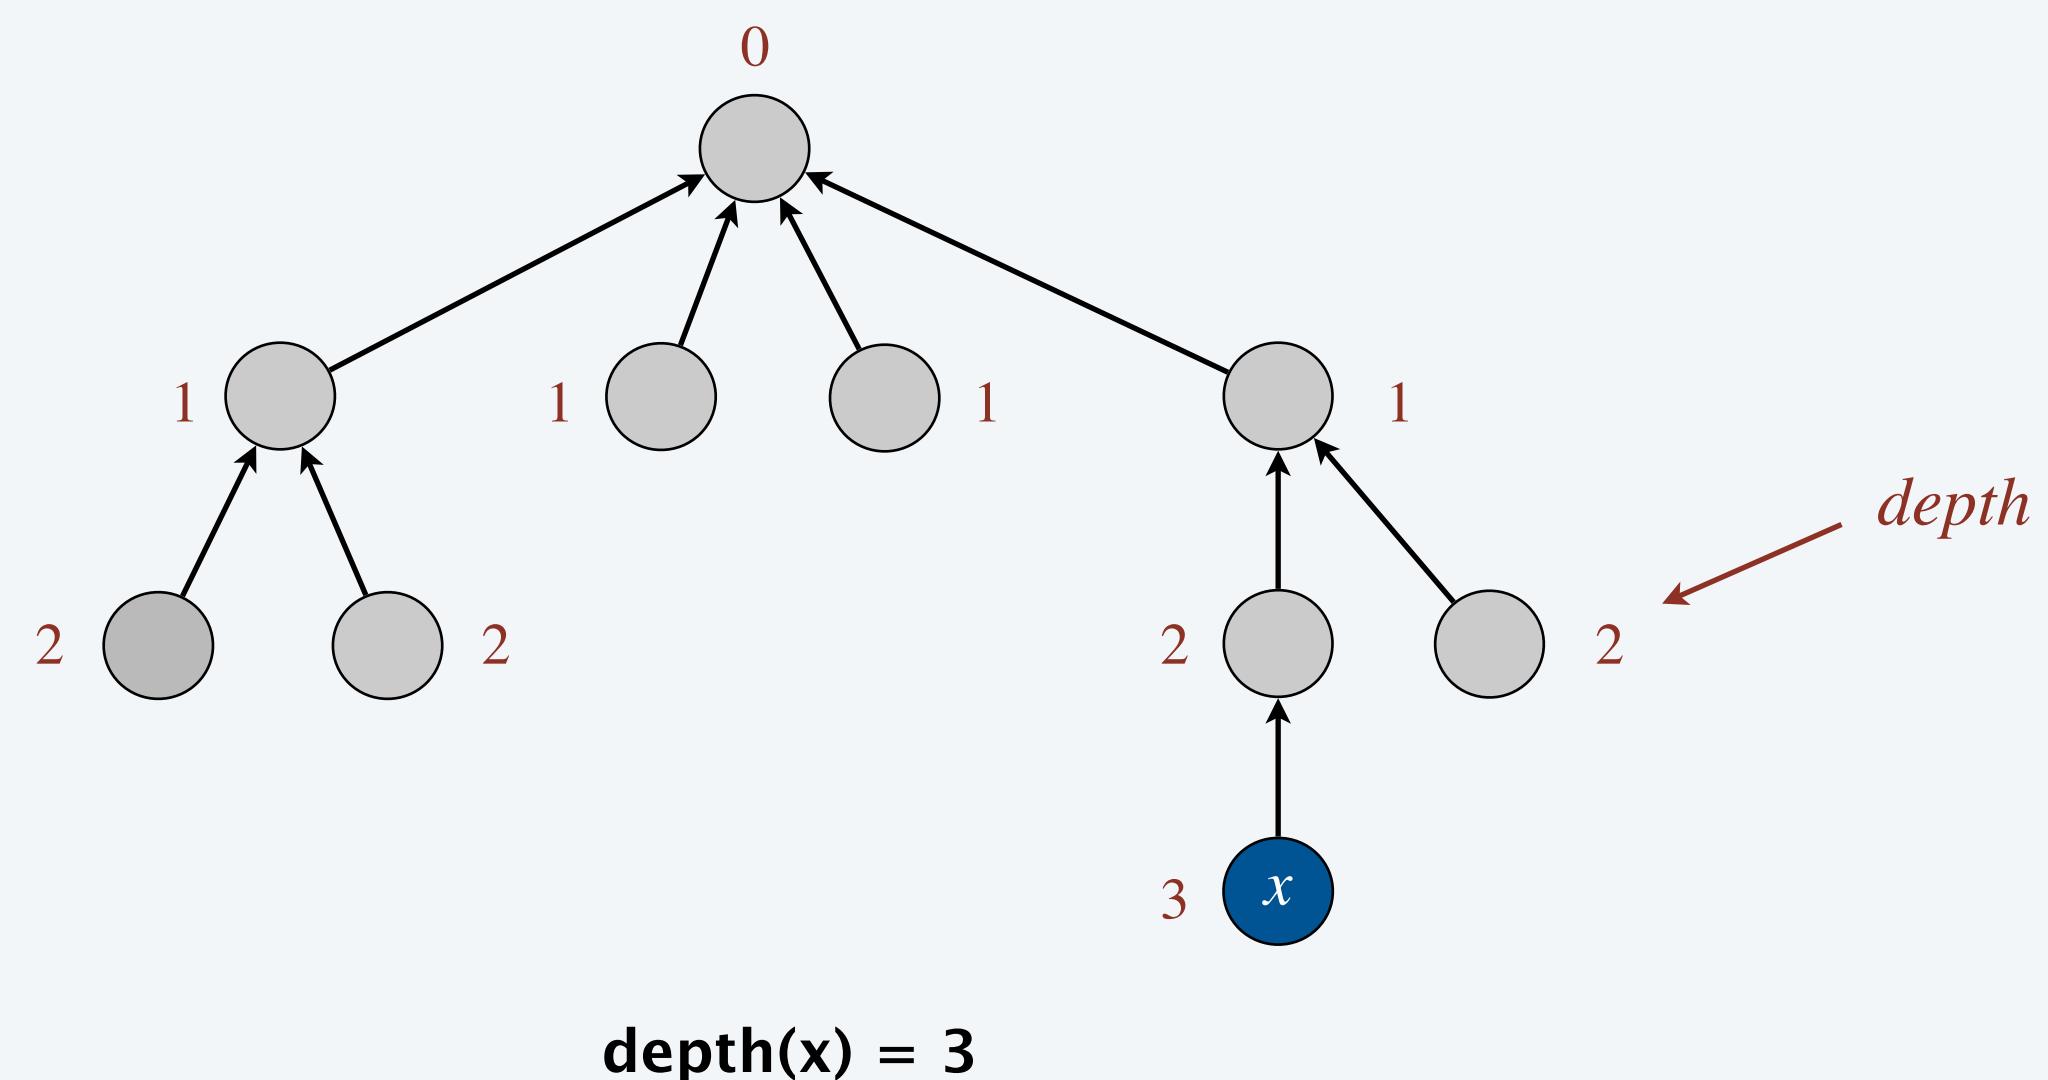
Quick-union analysis

Cost model. Number of array accesses (for read or write).

Running time.

- `union()` takes constant time, given two roots.
- `find()` takes time proportional to **depth** of node in tree.

*# links on path
from node to root*



Quick-union analysis

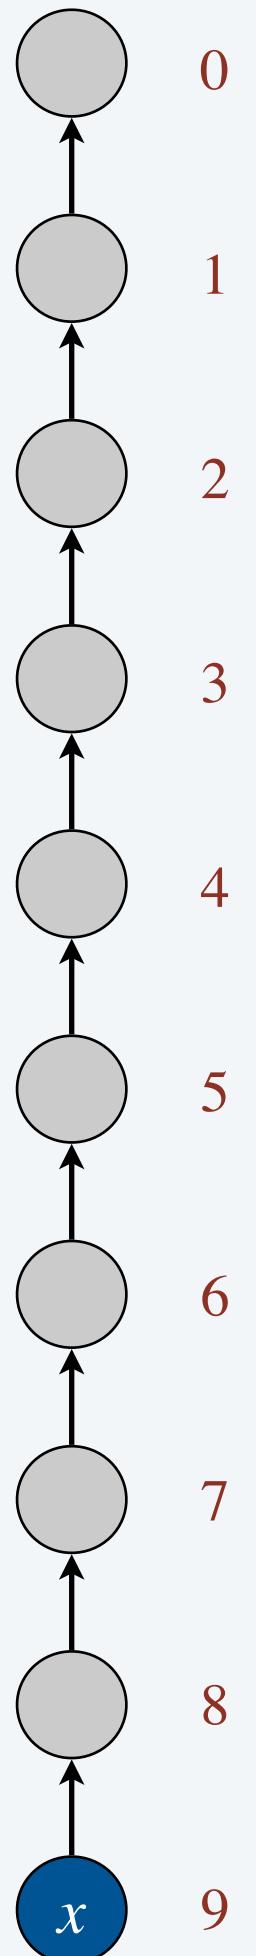
Cost model. Number of array accesses (for read or write).

Running time.

- `union()` takes constant time, given two roots.
- `find()` takes time proportional to **depth** of node in tree.

algorithm	initialize	union	find
quick-find	n	n	1
quick-union	n	n	n

worst-case number of array accesses (ignoring leading coefficient)



Union and find are too expensive (if trees get tall). Processing some sequences of m `union()` and `find()` operations on n elements takes $\geq mn$ array accesses.



quadratic in input size !

worst-case depth = $n-1$



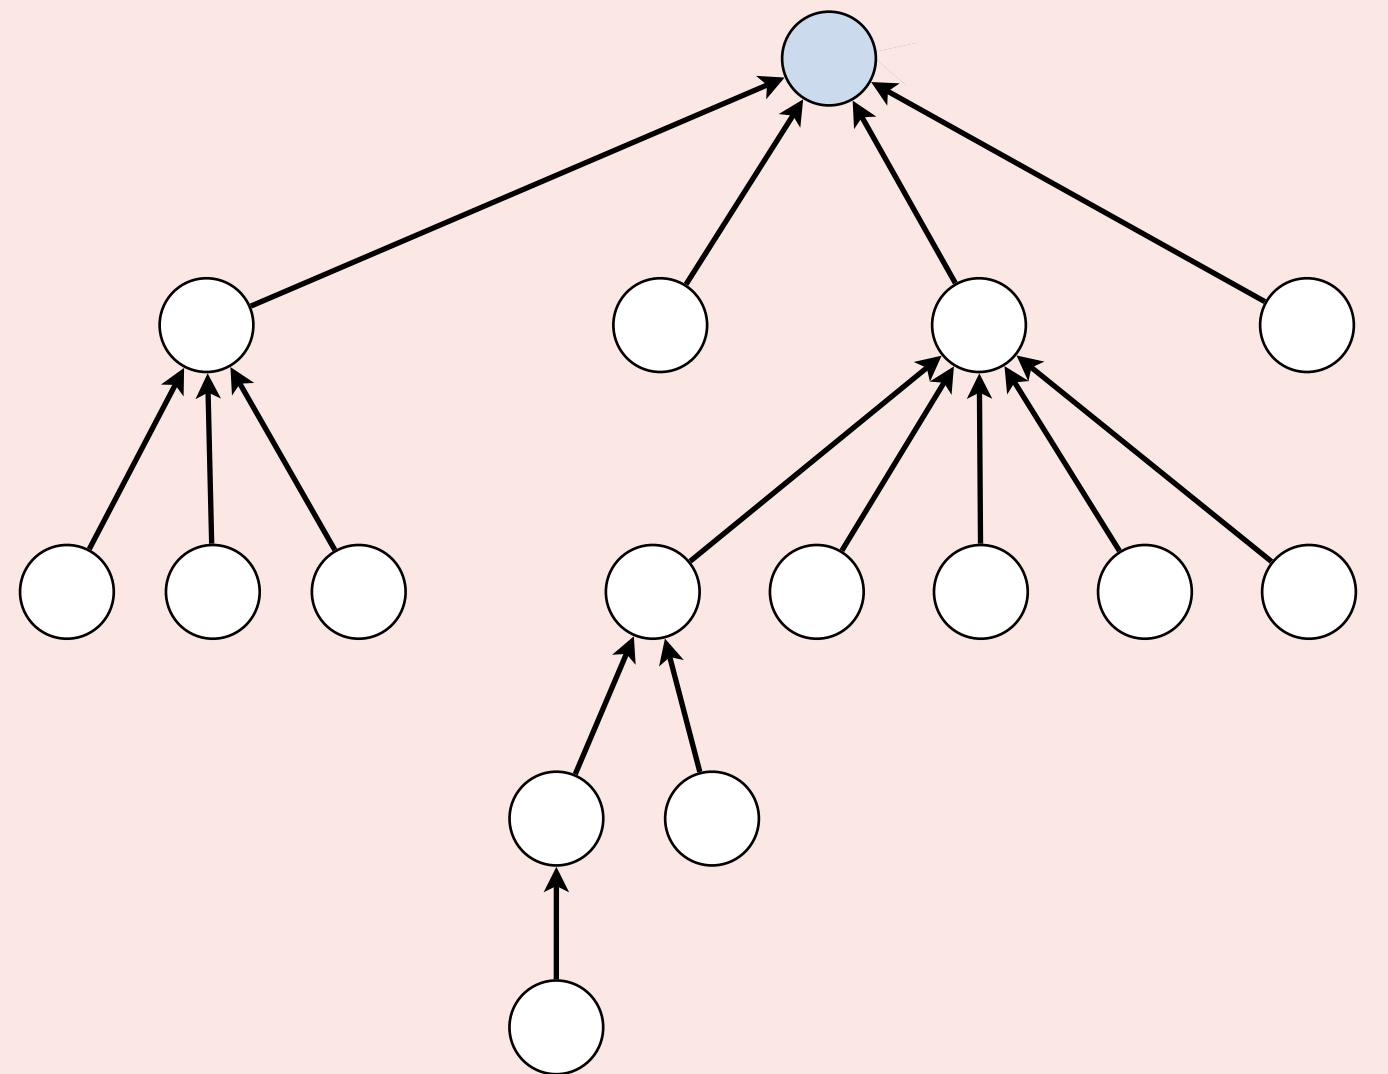
1.5 UNION-FIND

- ▶ *union-find data type*
- ▶ *quick-find*
- ▶ *quick-union*
- ▶ ***weighted quick-union***

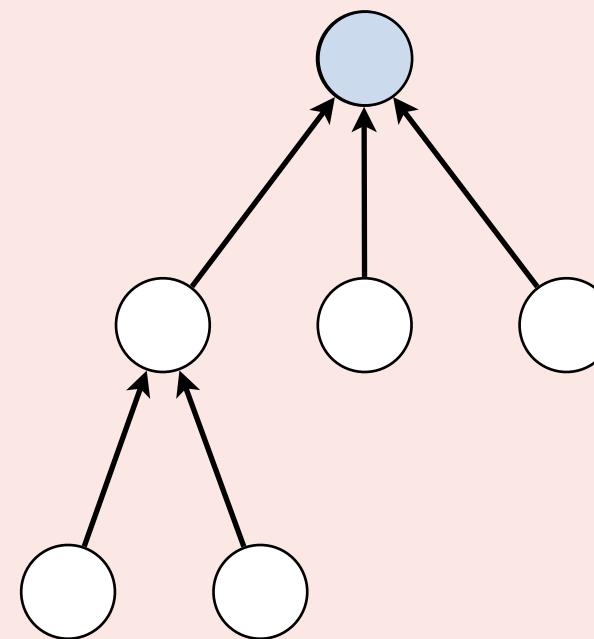


When linking two trees, which of these strategies is most effective?

- A. Link the root of the **smaller** tree to the root of the **larger** tree.
- B. Link the root of the **larger** tree to the root of the **smaller** tree.
- C. Flip a coin; randomly choose between A and B.
- D. All of the above.



larger tree
(size = 16, height = 4)

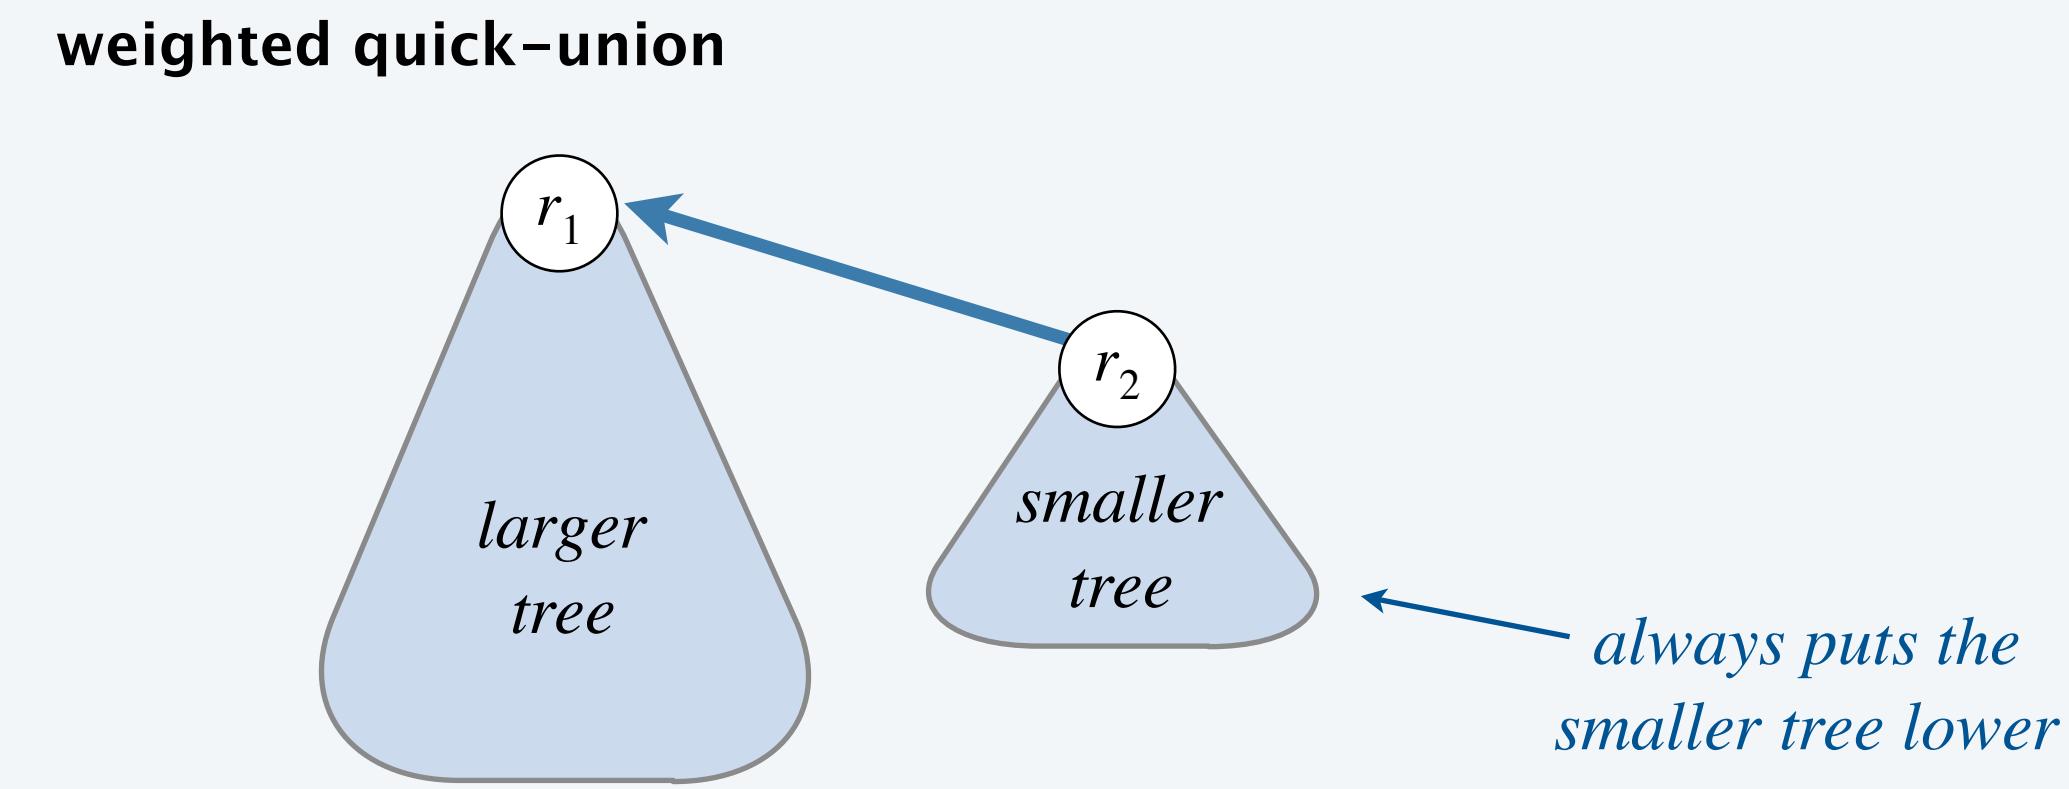
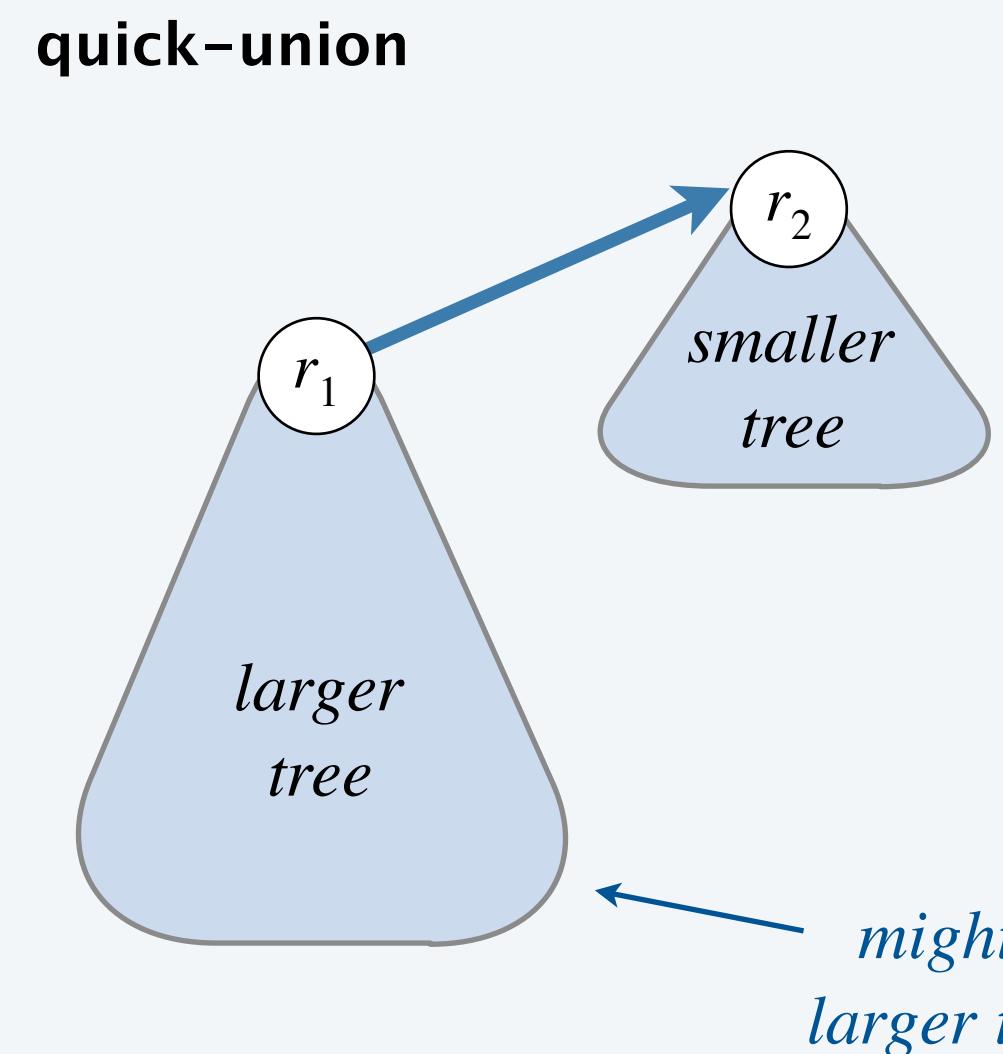


smaller tree
(size = 6, height = 2)

Weighted quick-union (link-by-size)

Link-by-size. Modify quick-union to avoid tall trees.

- Keep track of **size** of each tree = number of elements.
- Always link root of smaller tree to root of larger tree. ← *fine alternative: link-by-height
(minimize worst-case depth vs. average depth)*



Weighted quick-union: Java implementation

Data structure. Same as quick-union, but maintain extra array `size[i]`

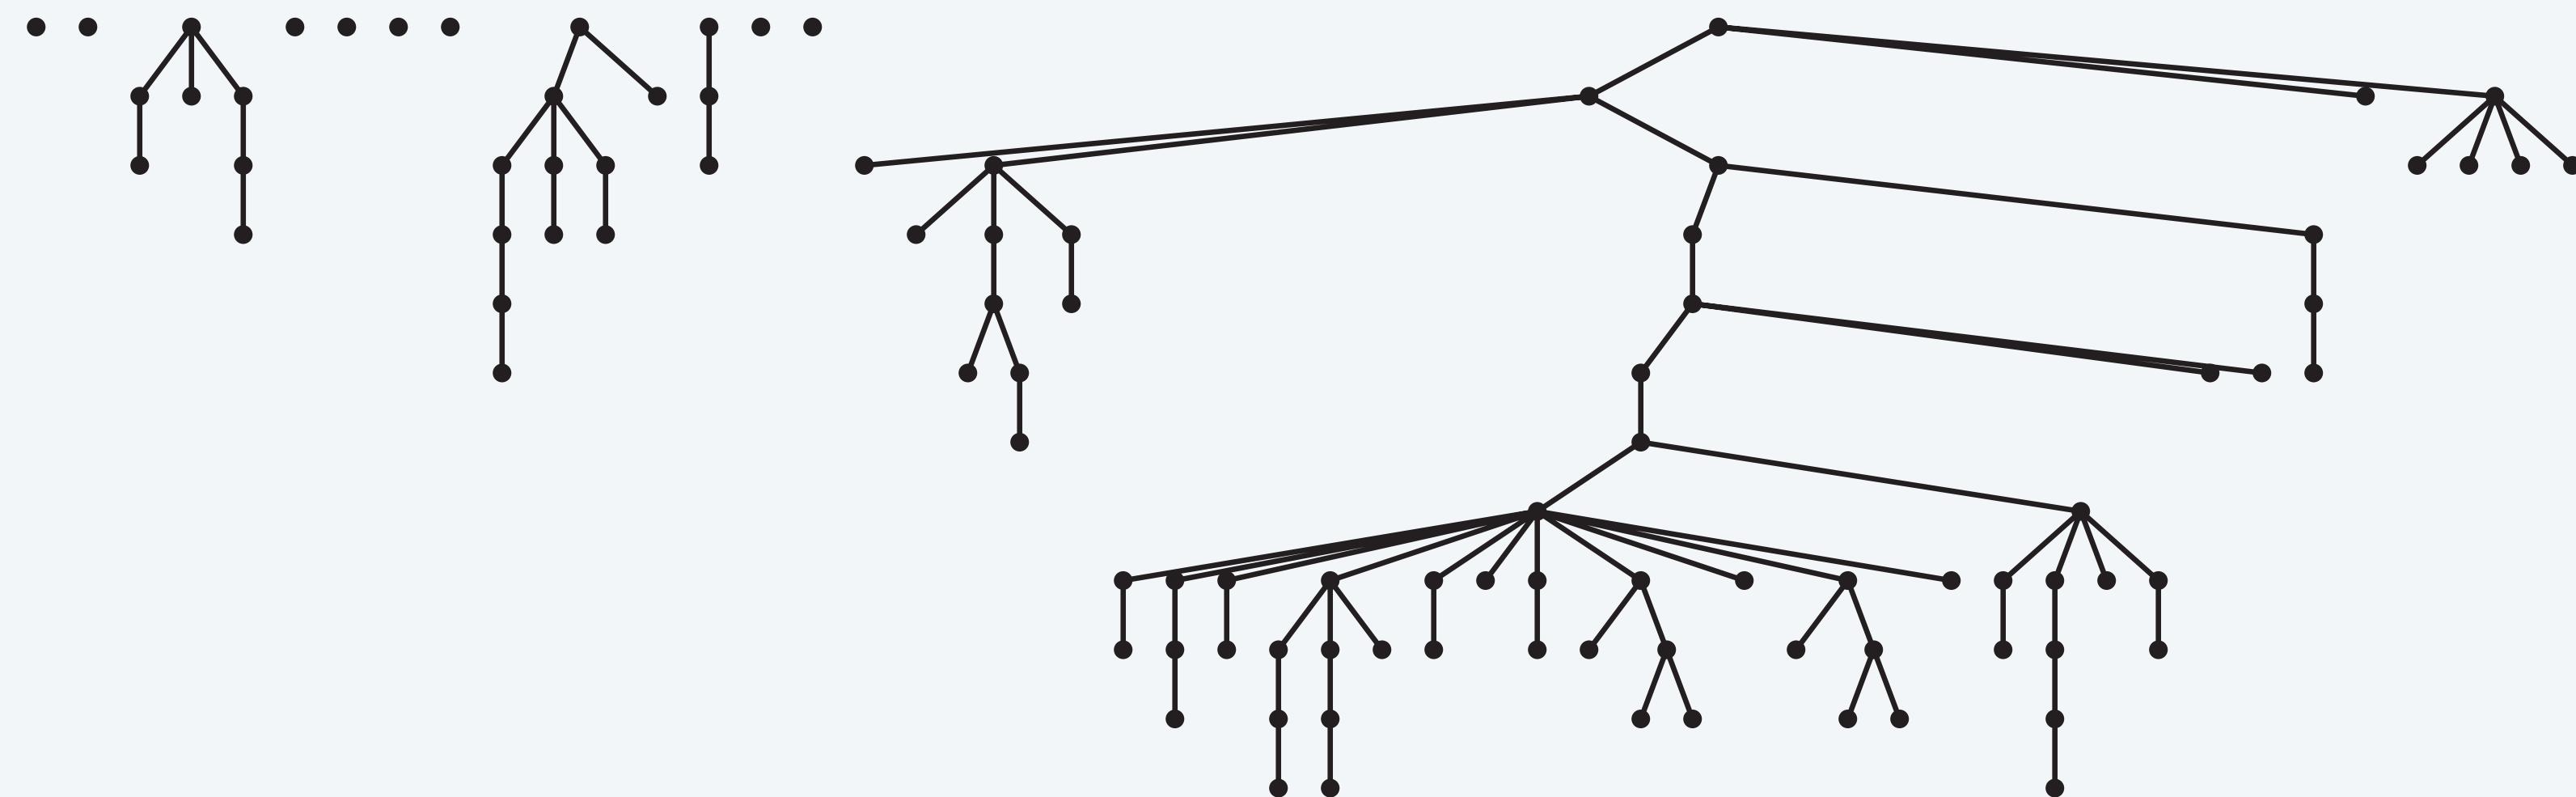
to count number of elements in the tree rooted at `i`, initially `1`.

- `find()`: identical to quick-union.
- `union()`: link root of smaller tree to root of larger tree; update `size[]`.

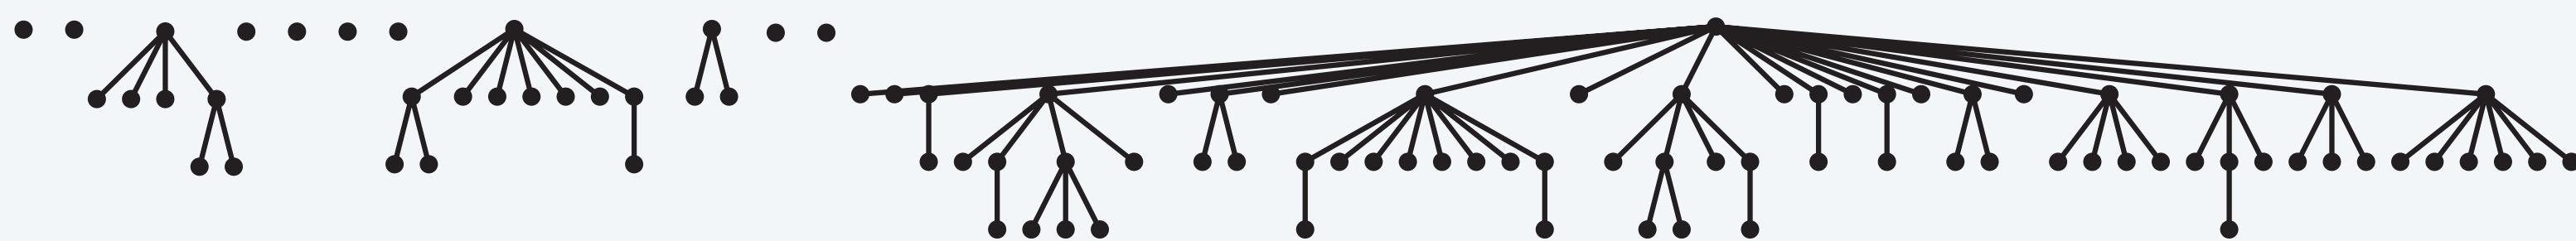
```
public void union(int p, int q) {  
    int rootP = find(p);  
    int rootQ = find(q);  
    if (rootP == rootQ) return; ← p and q already in the same set  
  
    if (size[rootP] < size[rootQ]) { ← link root of smaller tree  
        parent[rootP] = rootQ; to root of larger tree  
        size[rootQ] += size[rootP]; (and update size)  
    }  
    else {  
        parent[rootQ] = rootP;  
        size[rootP] += size[rootQ];  
    }  
}
```

Quick-union vs. weighted quick-union: larger example

quick-union

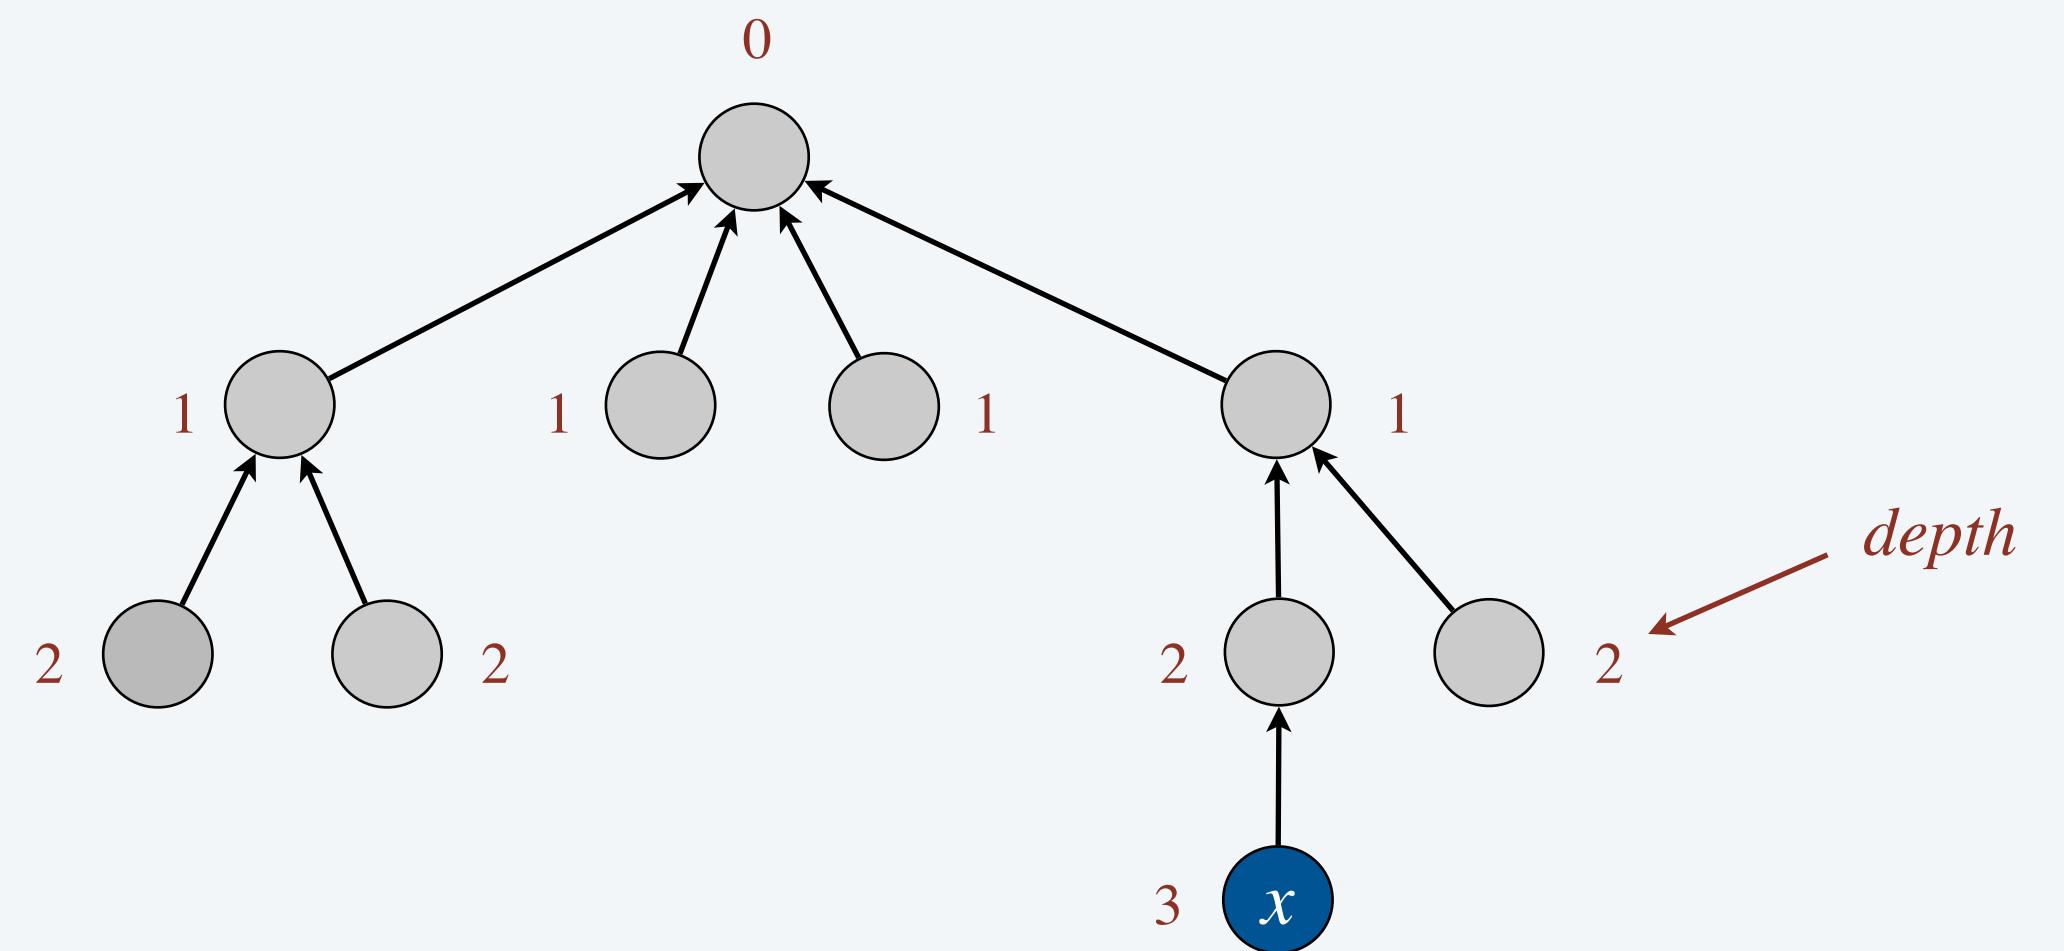


weighted



Weighted quick-union analysis

Proposition. Depth of any node $x \leq \log_2 n$.

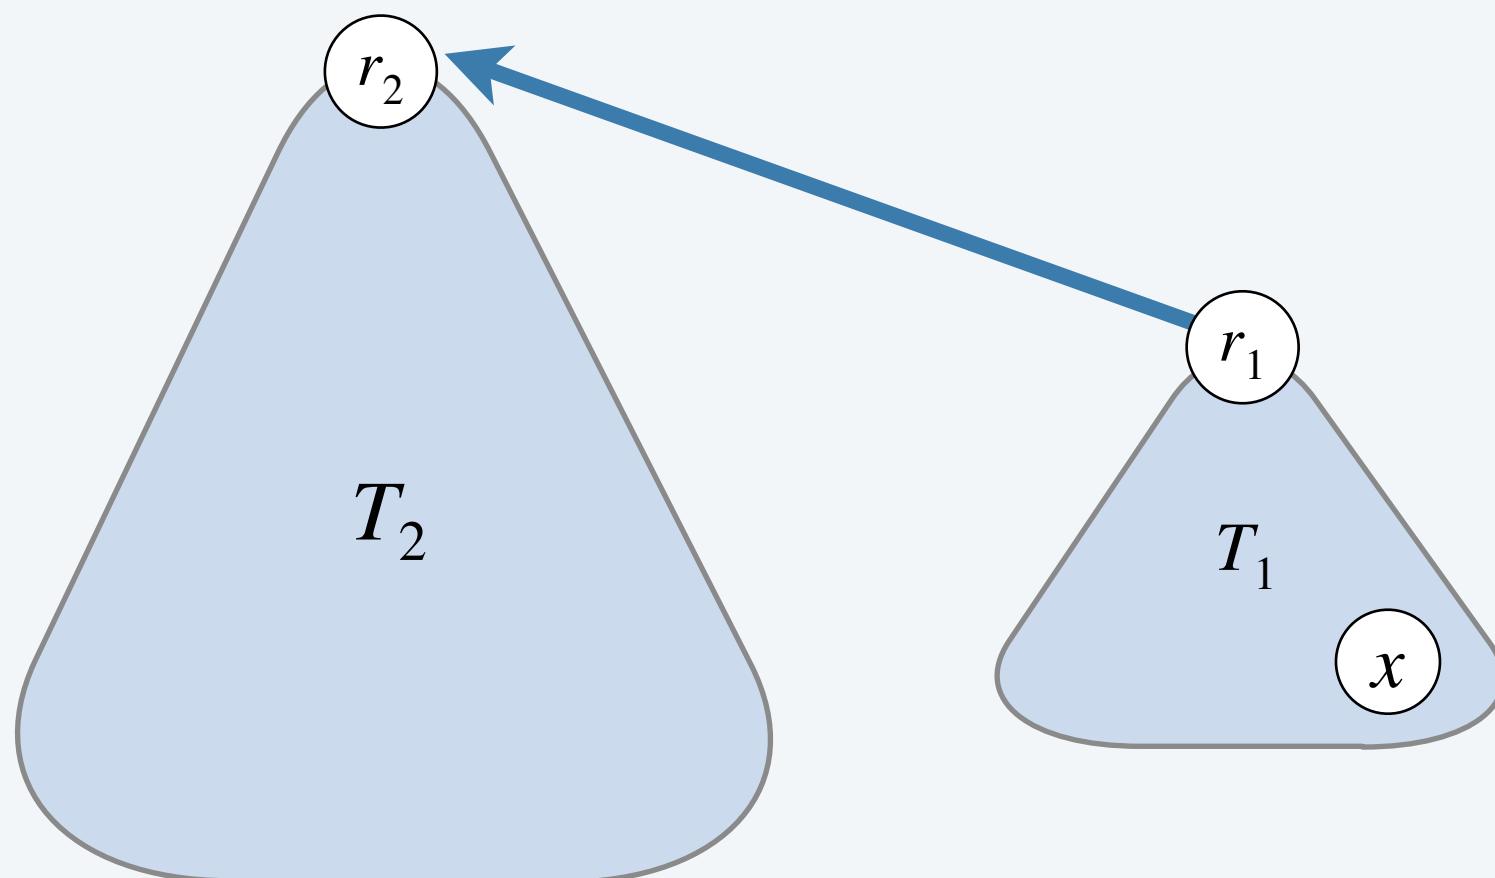
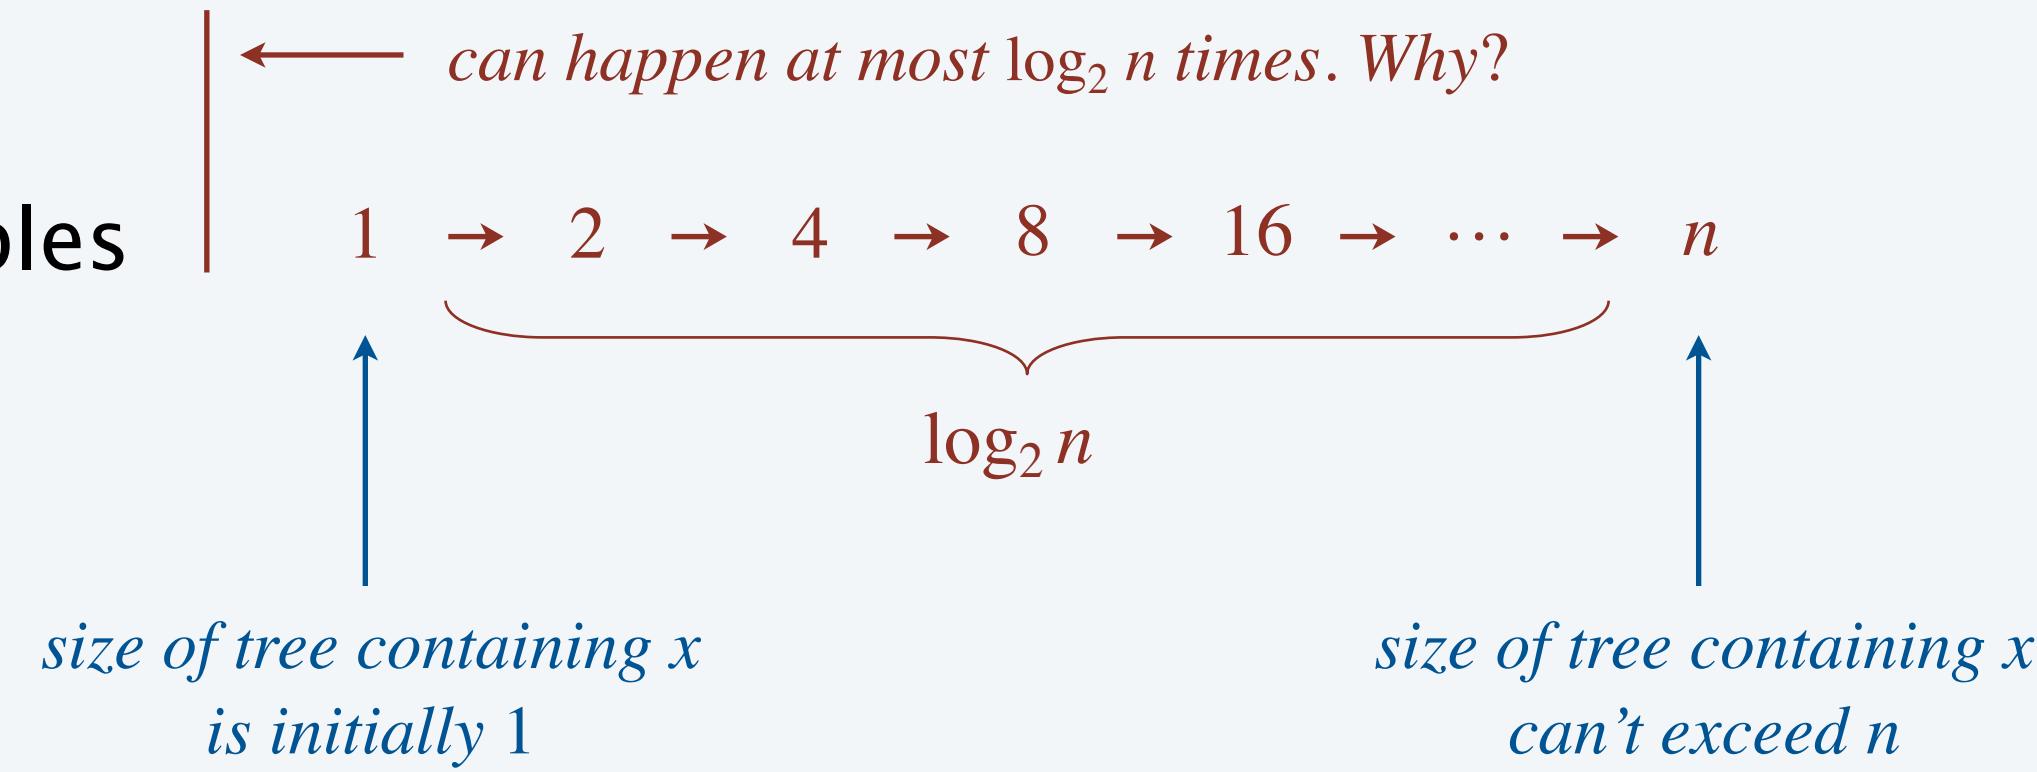


Weighted quick-union analysis

Proposition. Depth of any node $x \leq \log_2 n$.

Pf.

- Depth of x does not change unless root of tree T_1 containing x is linked to the root of a larger tree T_2 , forming a new tree T_3 .
- When this happens:
 - depth of x increases by exactly 1
 - size of tree containing x at least doubles because $\text{size}(T_3) = \text{size}(T_1) + \text{size}(T_2) \geq 2 \times \text{size}(T_1)$.



Weighted quick-union analysis

Proposition. Depth of any node $x \leq \log_2 n$.

Running time.

- `union()` takes constant time, given two roots.
- `find()` takes time proportional to **depth** of node in tree.

algorithm	initialize	union	find
quick-find	n	n	1
quick-union	n	n	n
weighted quick-union	n	$\log n$	$\log n$

worst-case number of array accesses (ignoring leading coefficient)

in this course, log mean logarithm for some constant base

Summary

Key point. Weighted quick-union empowers us to solve problems that could not otherwise be addressed.

algorithm	worst-case time
quick-find	$m n$
quick-union	$m n$
weighted quick-union	$m \log n$
quick-union + path compression	$m \log n$
weighted quick-union + path compression	$m \alpha(m, n)$

order of growth for $m \geq n$ union-find operations on a set of n elements

Ex. [10^9 union-find operations on 10^9 elements]

- Efficient algorithm reduces time from 30 years to 6 seconds.
- Supercomputer won't help much.

Credits

image	source
<i>Game of Hex</i>	Wolfram MathWorld
<i>Cluster Labeling</i>	Tiberiu Marita
<i>Bob Tarjan</i>	Princeton University
<i>Computer and Supercomputer</i>	New York Times

A final thought

*“The goal is to come up with algorithms that you can apply in practice that **run fast**, as well as being **simple, beautiful, and analyzable**.”* — Robert Tarjan

