Splay Trees

Outline for Today

- ***Static Optimality***
  - Balanced BSTs aren't necessarily optimal!
- ***Splay Trees***
  - A self-adjusting binary search tree.
- ***Properties of Splay Trees***
  - Why is splaying worthwhile?
- ***Dynamic Optimality (ITA)***
  - An open problem in data structures.

Static Optimality

Balanced BSTs

- Balanced BSTs guarantee tree operations run in worst-case time O(log *n*).
- ***Claim:*** If the elements in the tree aren't accessed uniformly, then a balanced BST might not actually be the ideal BST.

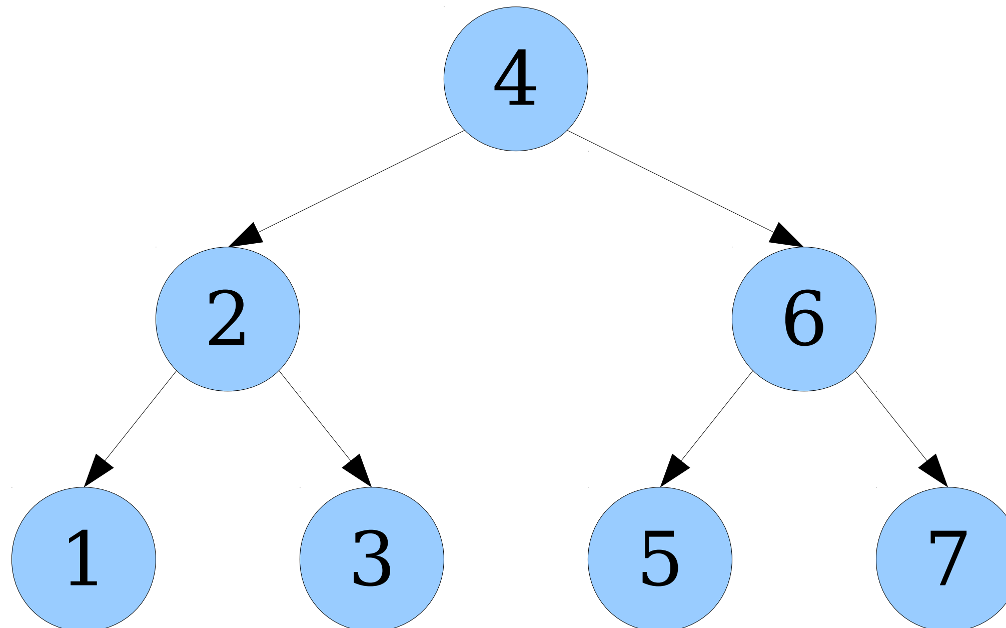

Balanced BSTs

- Balanced BSTs guarantee tree operations run in worst-case time O(log *n*).
- ***Claim:*** If the elements in the tree aren't accessed uniformly, then a balanced BST might not actually be the ideal BST.

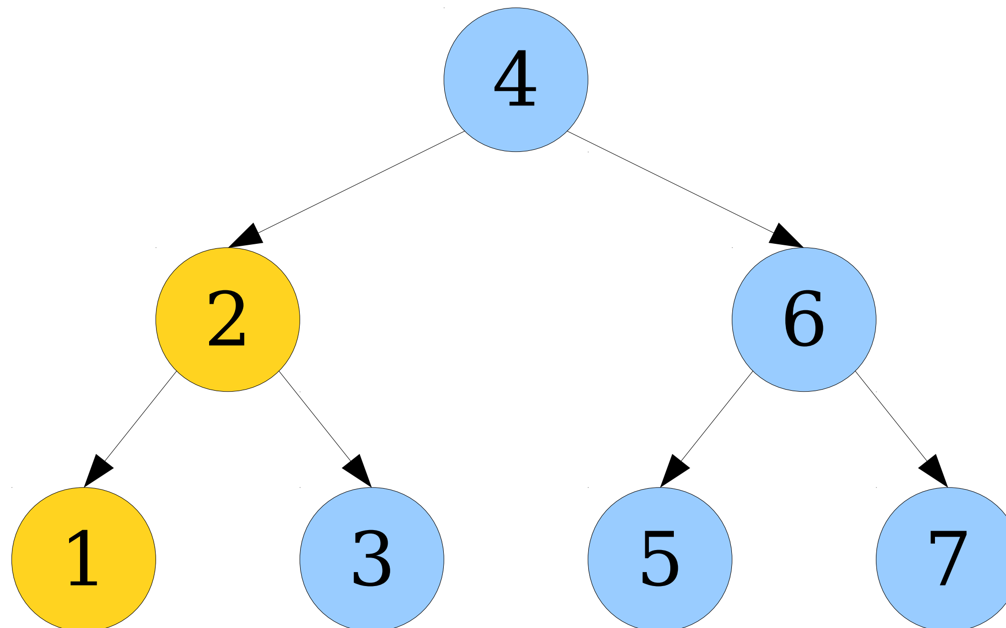

Balanced BSTs

- Balanced BSTs guarantee tree operations run in worst-case time O(log *n*).
- ***Claim:*** If the elements in the tree aren't accessed uniformly, then a balanced BST might not actually be the ideal BST.

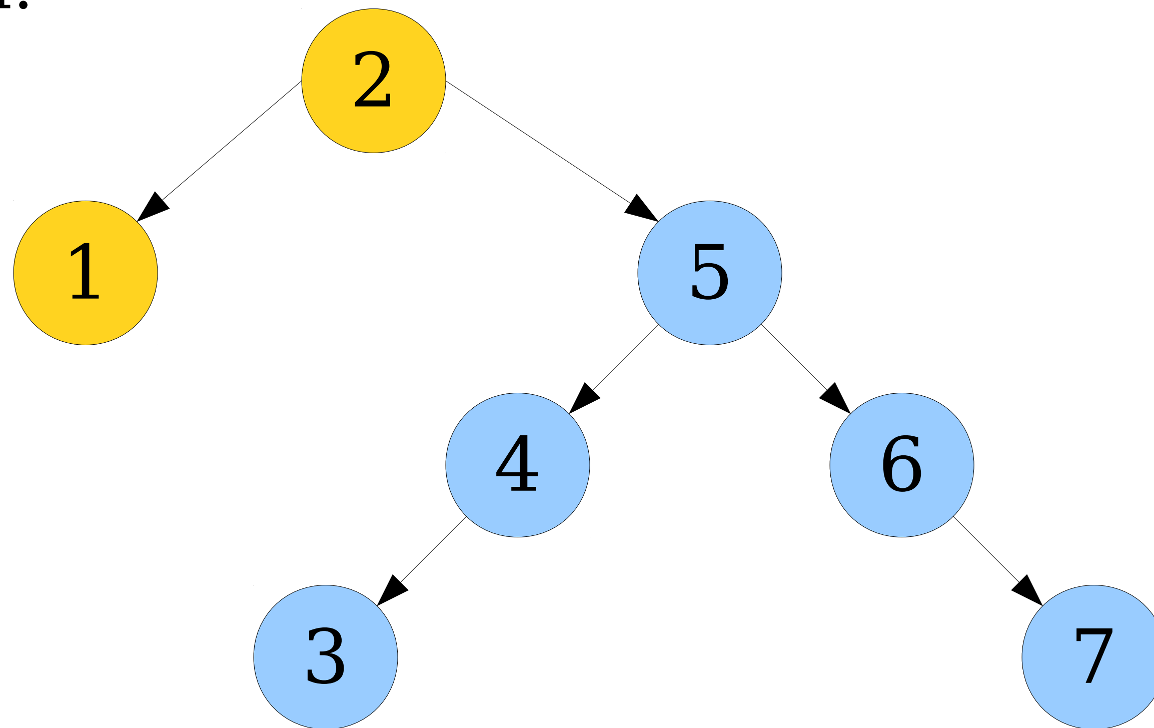

Balanced BSTs

- Balanced BSTs guarantee tree operations run in worst-case time O(log n).
- ***Claim:*** If the elements in the tree aren't accessed uniformly, then a balanced BST might not actually be the ideal BST.

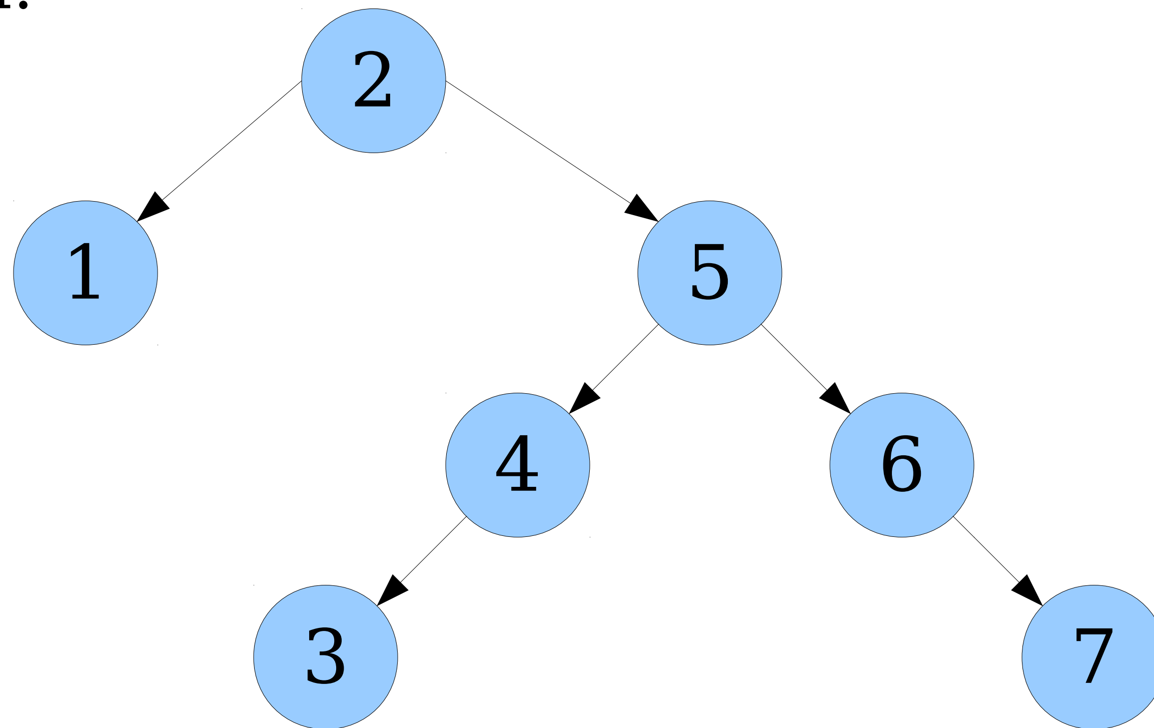

Balanced BSTs

- Balanced BSTs guarantee tree operations run in worst-case time O(log *n*).
- ***Claim:*** If the elements in the tree aren't accessed uniformly, then a balanced BST might not actually be the ideal BST.

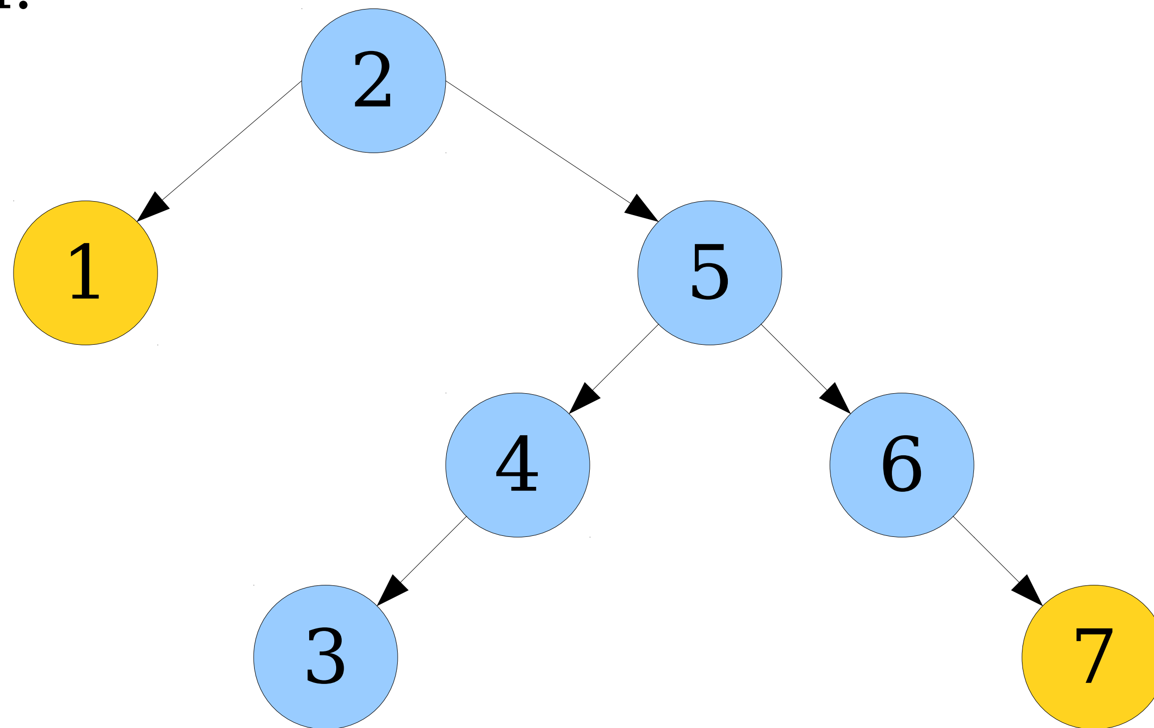

Balanced BSTs

- Balanced BSTs guarantee tree operations run in worst-case time O(log n).
- ***Claim:*** If the elements in the tree aren't accessed uniformly, then a balanced BST might not actually be the ideal BST.

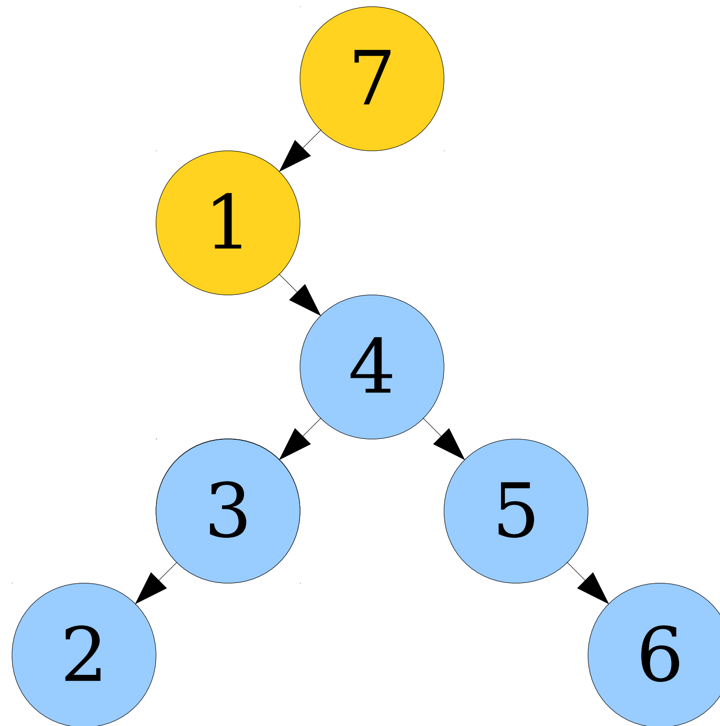

Balanced BSTs

- Balanced BSTs guarantee tree operations run in worst-case time O(log *n*).
- ***Claim:*** If the elements in the tree aren't accessed uniformly, then a balanced BST might not actually be the ideal BST.

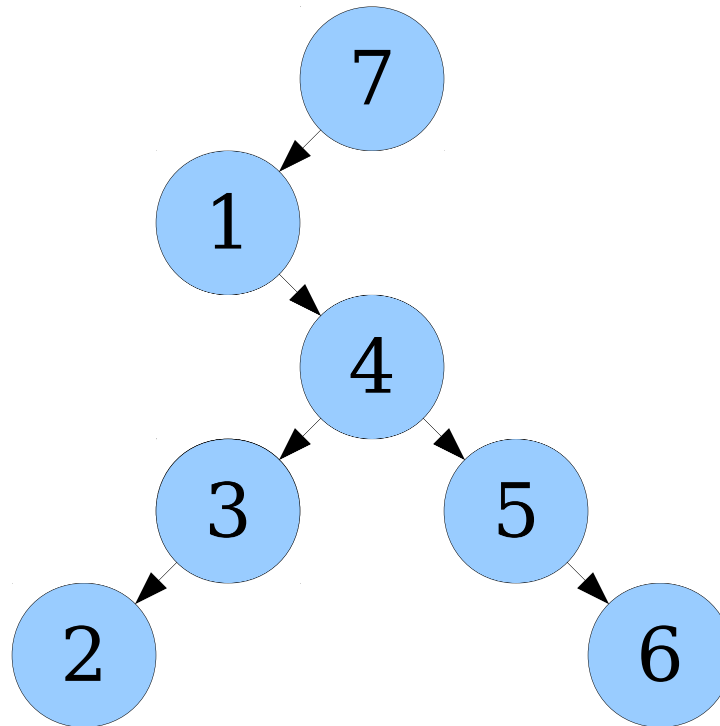

Static Optimality

- Let $S = \{ x_1, x_2, \ldots, x_n \}$ be a set with access probabilities $p_1, p_2, \ldots, p_n$.
- If T is a BST whose keys are the keys in S, then let X_T be a random variable equal to the number of nodes in T that are touched when performing a lookup, assuming the key to look up is sampled from the above probability distribution.
- ***Goal:*** Construct a binary search tree T^* such that $E[X_{T^*}]$ is minimal.
- T^* is called a ***statically optimal binary search tree***.

Static Optimality

- ***Theorem:*** There is an O(n^2)-time dynamic programming algorithm for constructing statically optimal binary search trees.
  - Knuth, 1971. (See CLRS)
- ***Theorem:*** Weight-balanced trees whose weights are the element access probabilities have an expected lookup cost with a factor of 1.5 of a statically-optimal tree.
  - Mehlhorn, 1975.
- You can build a weight-balanced tree for a set of keys in time O(n) using a clever divide-and-conquer algorithm. You'll see this in PS4.

Finding a Lower Bound

- Suppose we design a BST data structure where the cost of any lookup is $O(\log n)$.
- You'd intuitively know that, at least from the perspective of worst-case efficiency, you couldn't improve on the cost of a lookup by more than a constant factor.
- Why is this?
- ***Justification:*** *Every* BST with n elements has a worst-case lookup time of $\Omega(\log n)$, since some element has to be at depth at least $\Omega(\log n)$.
- An $O(\log n)$ upper bound matches this lower bound, and therefore can't be improved.

Finding a Lower Bound

- Right now, we have an $\Omega(\log n)$ lower bound on the *worst-case* cost of a lookup in a BST.
- ***Question:*** Can we find some sort of lower bound on the *expected* cost of a lookup in a BST?
  - Here, the expectation is taken over the access probabilities.

Static Optimality

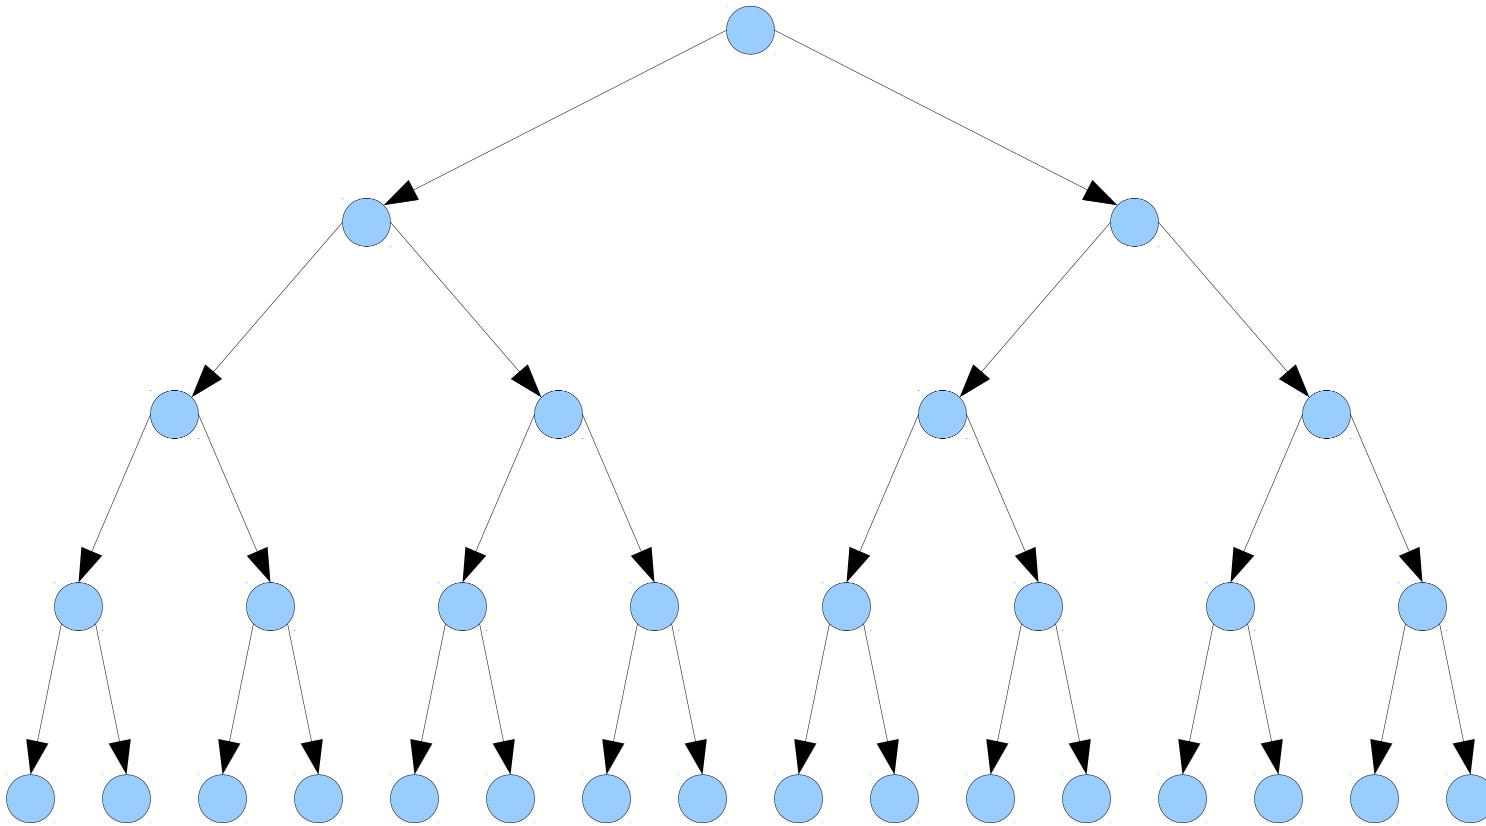

Static Optimality

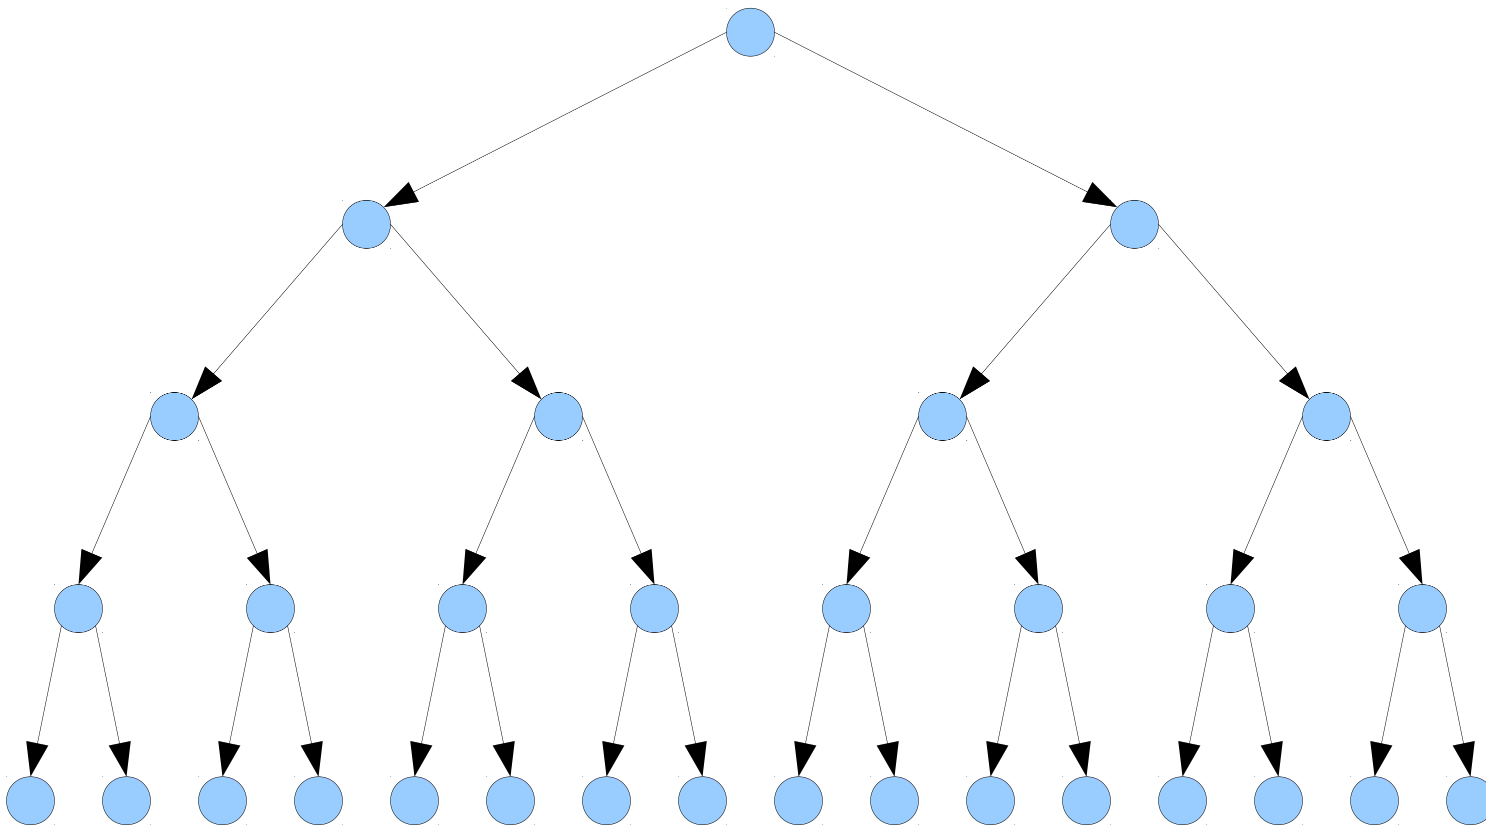

Intuition: Try to place nodes with higher access probabilities higher up in the tree.

Static Optimality

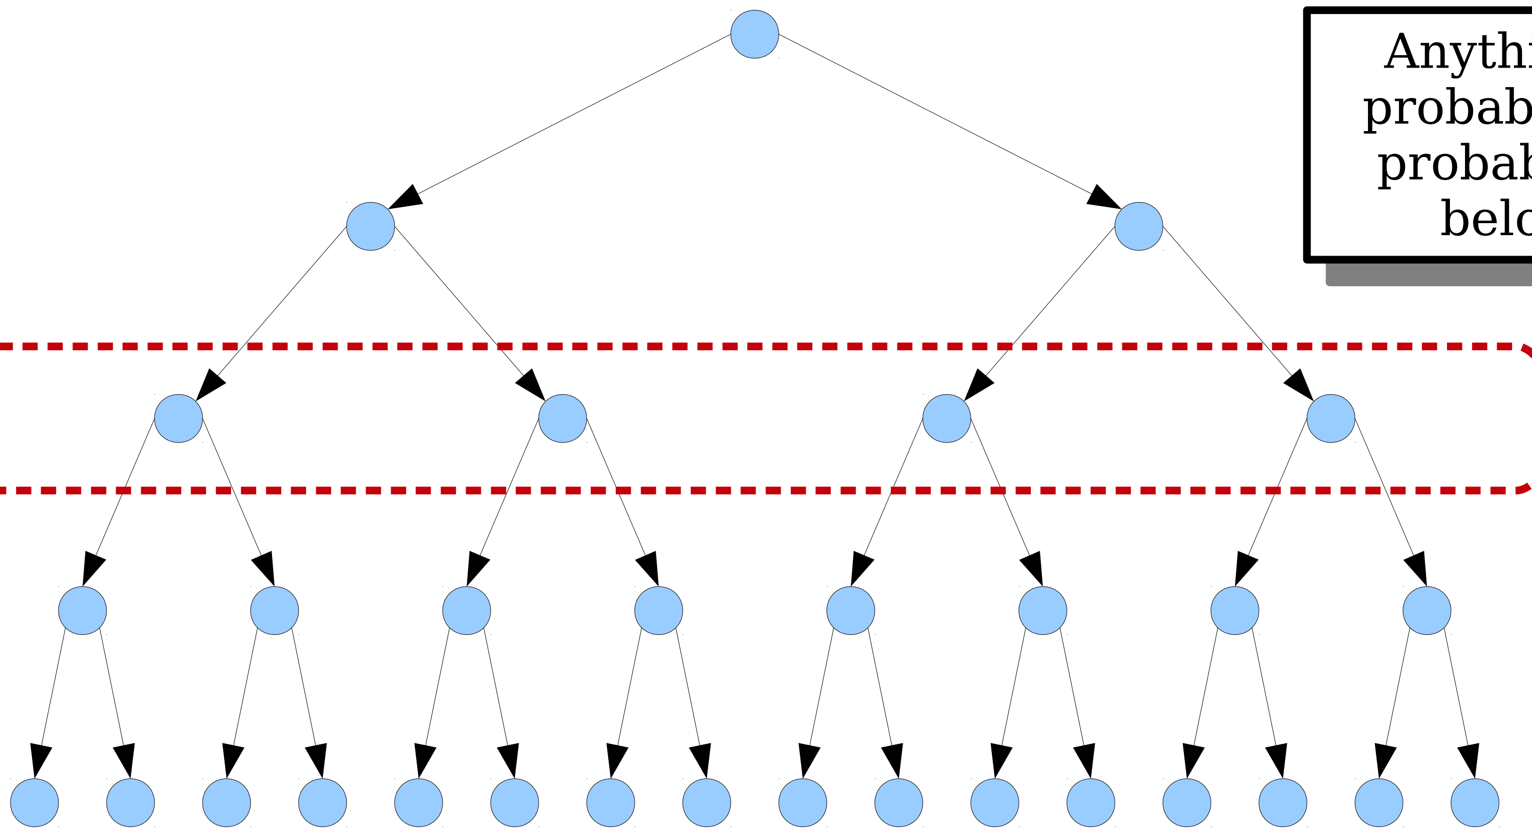

Anything with access probability 1/4 or above probably shouldn't be below this point.

Static Optimality

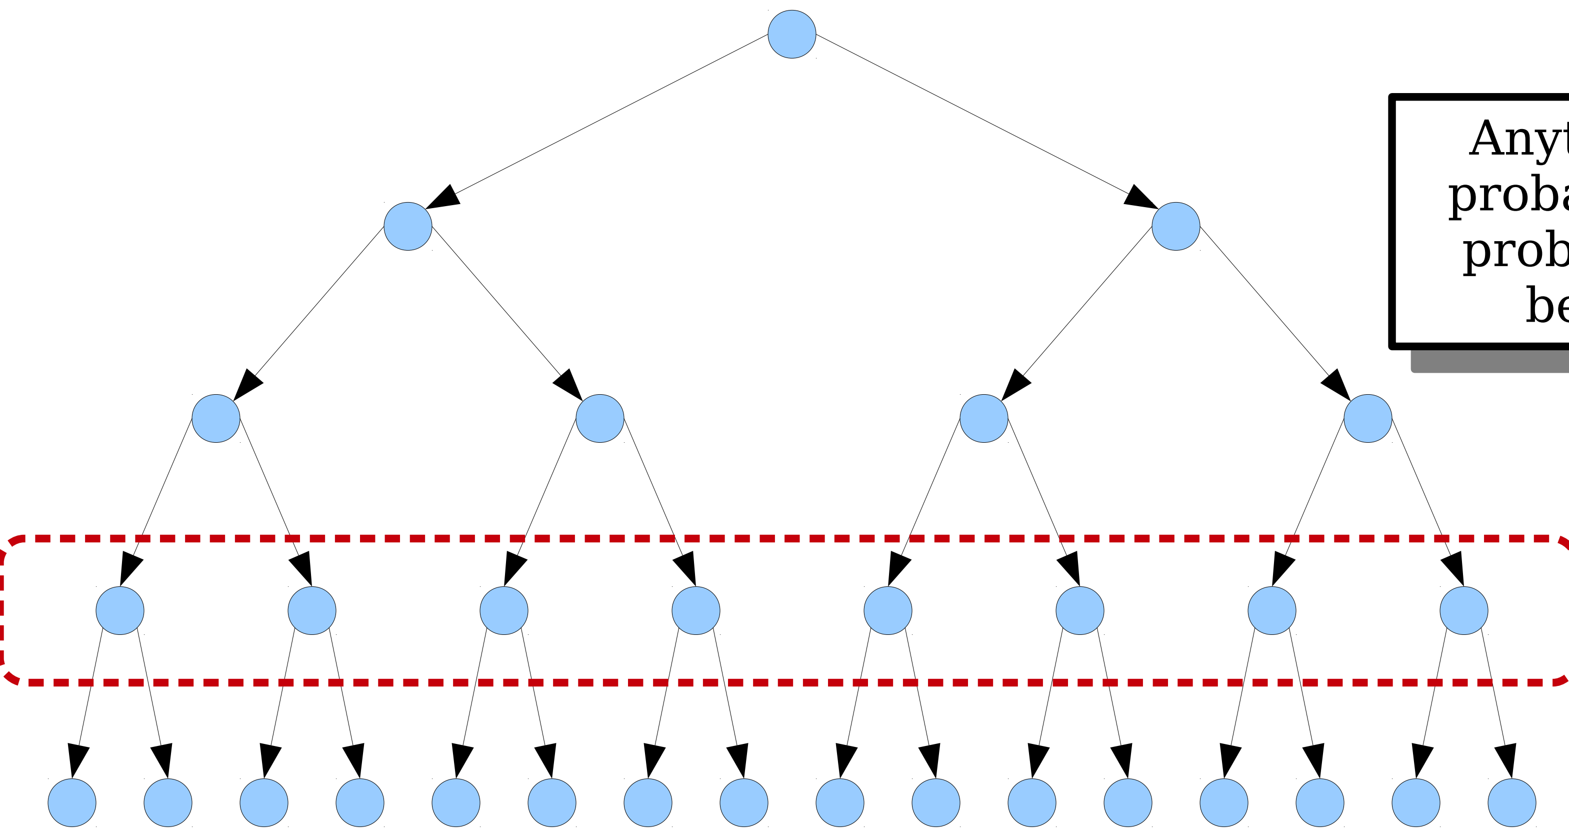

Anything with access probability $^{1}/_{8}$ or above probably shouldn’t be below this point.

Static Optimality

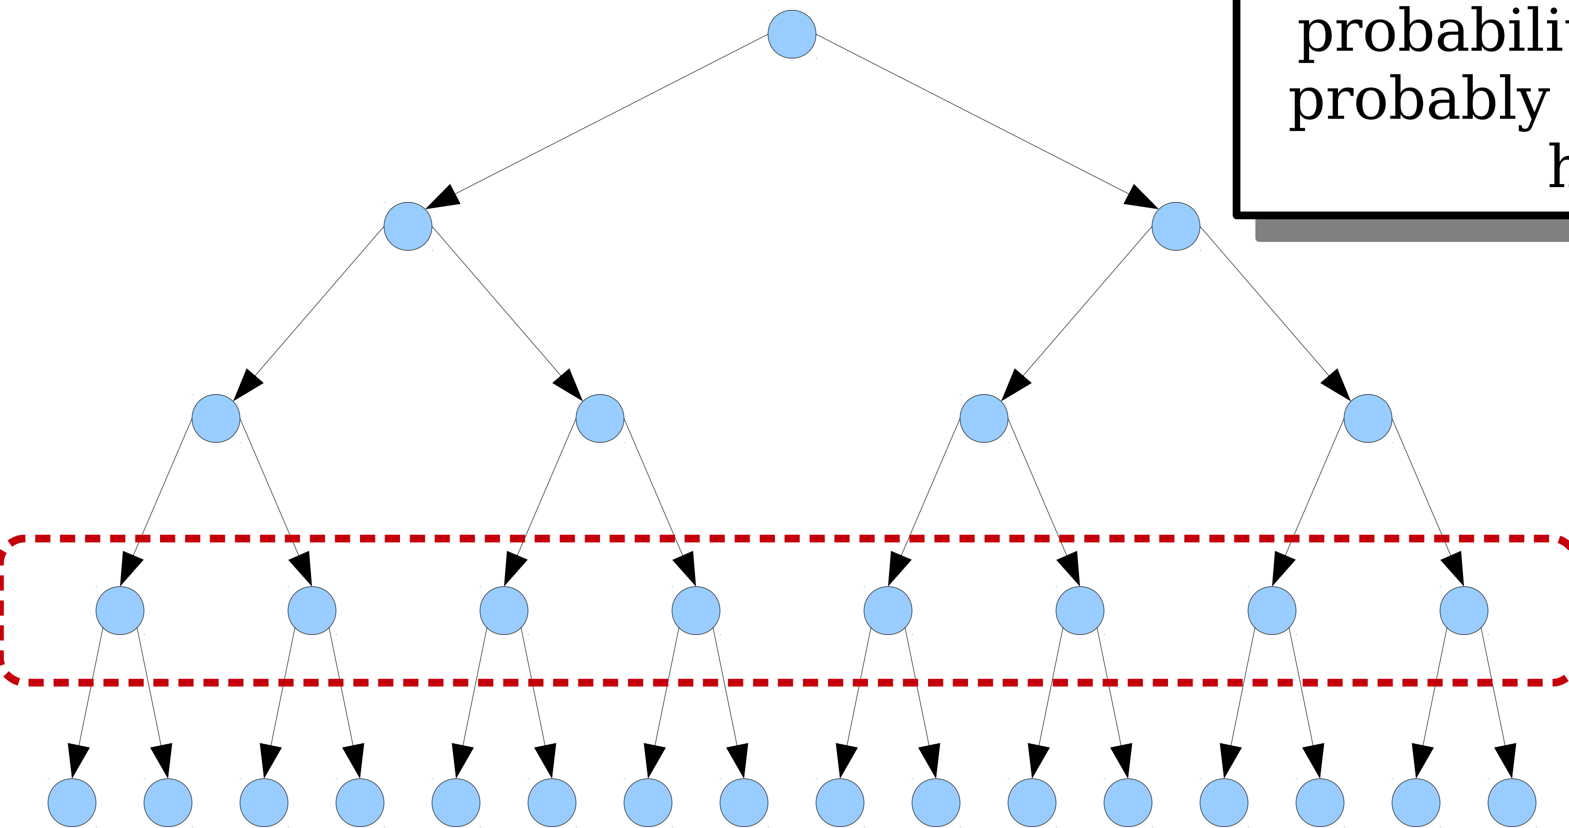

Idea: If a key has access probability 2^{-k}, it should probably go at level k or higher.

Static Optimality

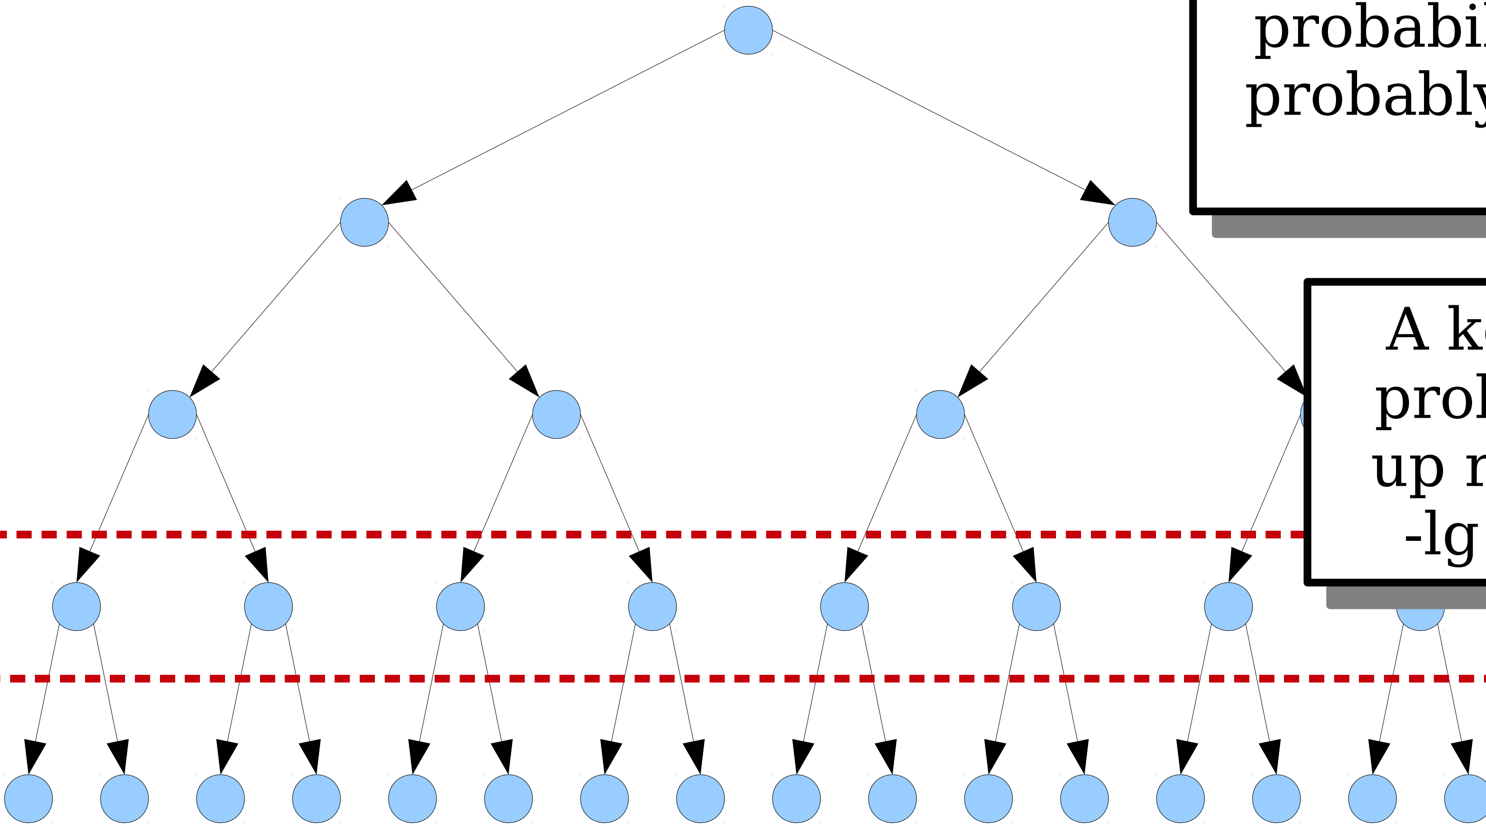

Idea: If a key has access probability 2^{-k}, it should probably go at level k or higher.

A key with access probability p_i ends up roughly at level $-\lg p_i$ in the tree.

Static Optimality

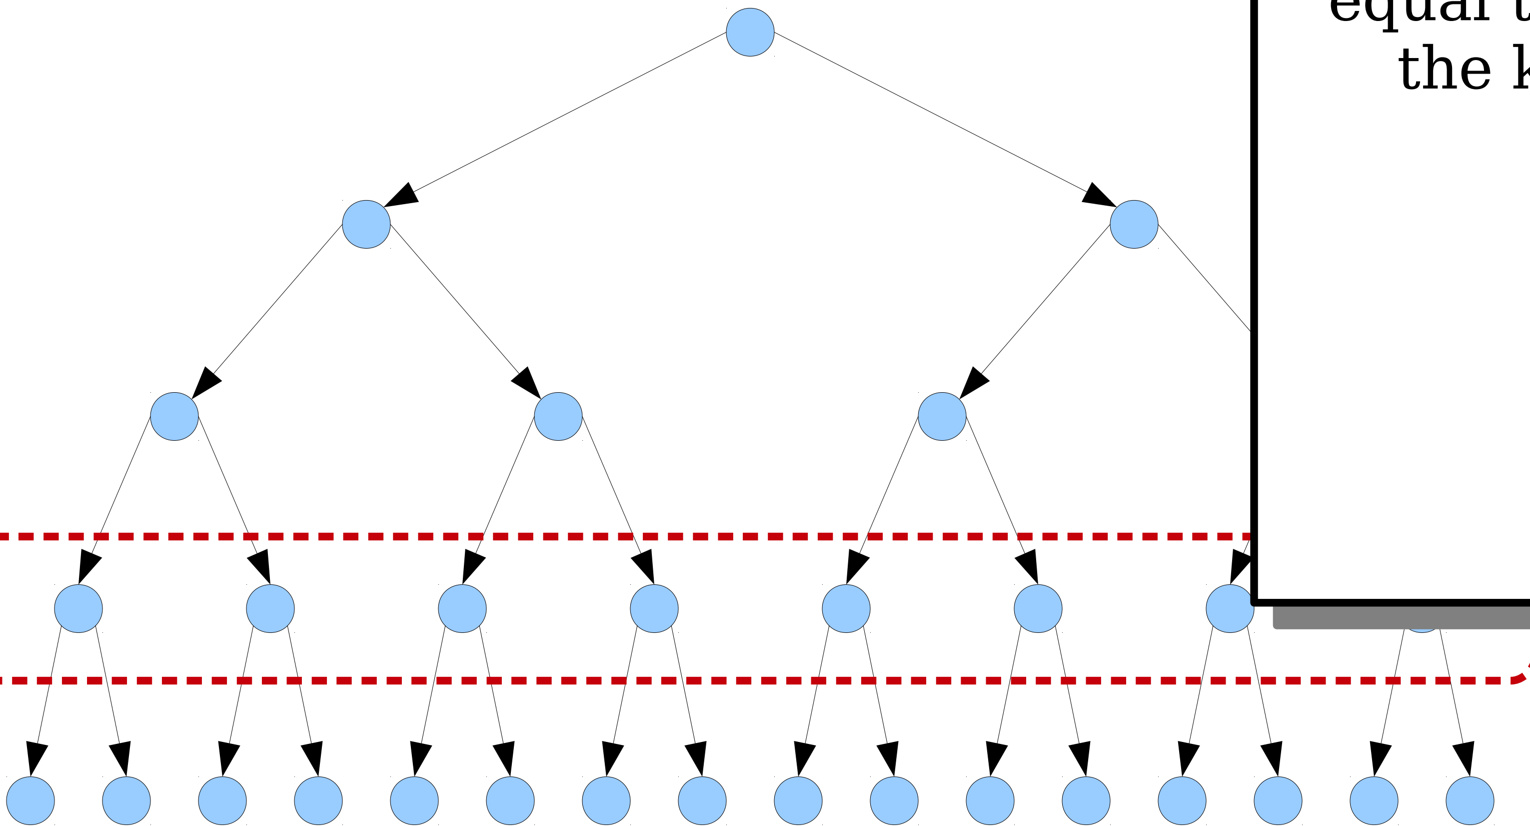

The cost of looking up some key x_i is roughly equal to the depth of the key x_i: $-\lg p_i$.

Static Optimality

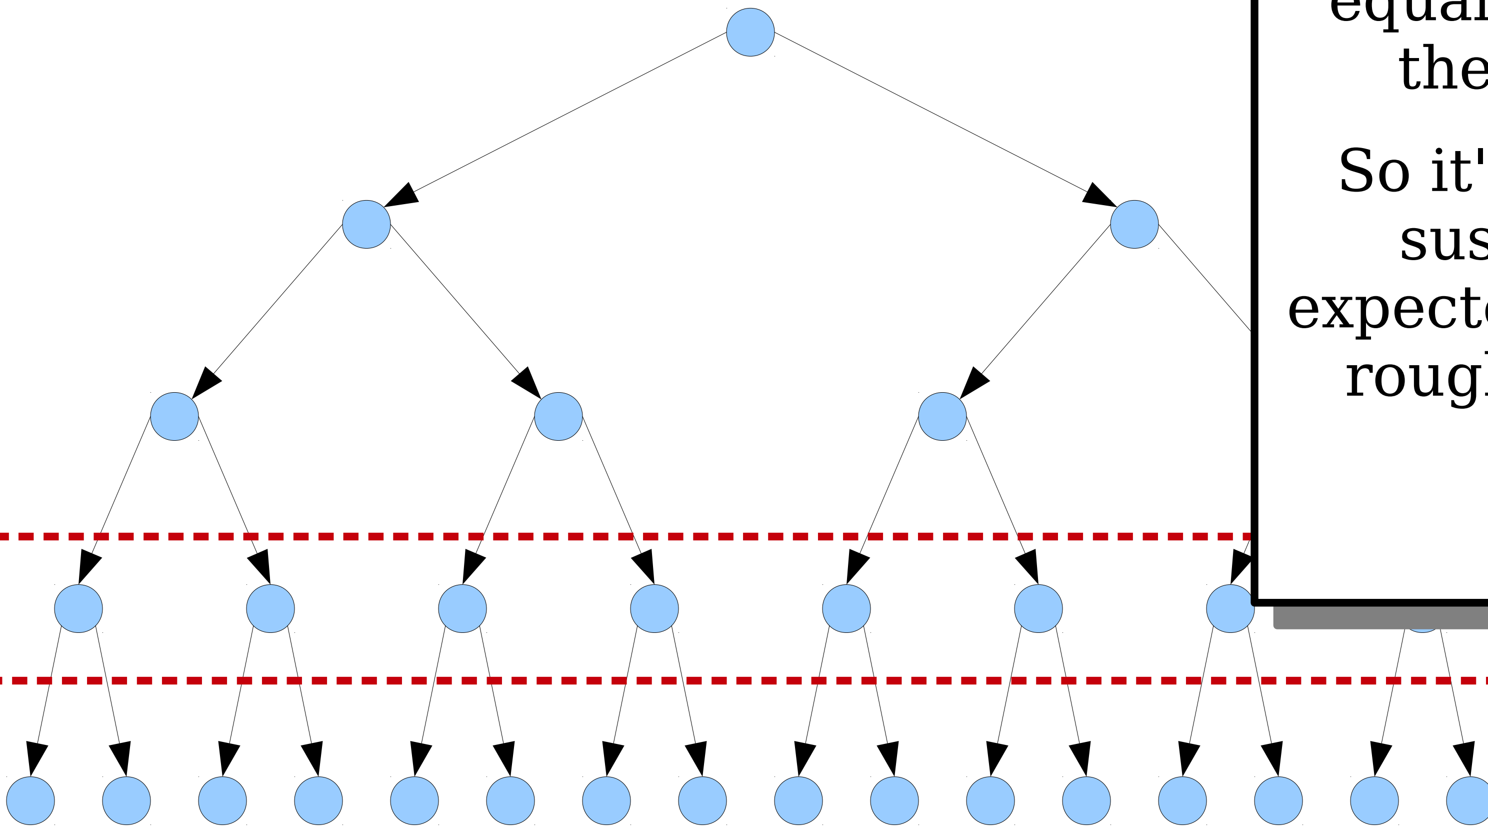

The cost of looking up some key x_i is roughly equal to the depth of the key x_i: -lg p_i.

So it's reasonable to suspect that the expected lookup cost to roughly work out to

$$\sum_{i=1}^{n} -p_i \lg p_i .$$

Shannon Entropy

- Consider a discrete probability distribution with elements $x_1, \ldots, x_n$, where element x_i has access probability p_i.
- The ***Shannon entropy*** of this probability distribution, denoted H_p (or just H, where p is implicit) is the quantity

$$H_p = \sum_{i=1}^{n} -p_i \lg p_i .$$

If we have n elements with equal access probability ($p_i = 1/n$), then

$$\begin{aligned} H_p &= \sum_{i=1}^{n} -p_i \lg p_i \\ &= \sum_{i=1}^{n} \frac{1}{n}\left(-\lg \frac{1}{n}\right) \\ &= \sum_{i=1}^{n} \frac{1}{n} \lg n \\ &= \mathbf{lg}\, \boldsymbol{n} \end{aligned}$$

If we have one element accessed 100% of the time ($p_1 = 1$) and all other elements are never accessed ($p_i = 0$), then

$$\begin{aligned} H_p &= \sum_{i=1}^{n} -p_i \lg p_i \\ &= -\lg 1 + \sum_{i=2}^{n} 0 \lg 0 \\ &= \mathbf{0} \end{aligned}$$

Static Optimality

- Consider a discrete probability distribution with elements $x_1, \ldots, x_n$, where element x_i has access probability p_i.
- The ***Shannon entropy*** of this probability distribution, denoted H_p (or just H, where p is implicit) is the quantity

$$H_p = \sum_{i=1}^{n} -p_i \lg p_i \, .$$

- ***Theorem:*** The expected lookup cost in *any* binary search tree for keys $x_1, \ldots, x_n$ with access probabilities $p_1, \ldots, p_n$ is $\Omega(1 + H)$.
- ***Theorem:*** For any set of keys x_i with access probabilities p_i, there is a BST that whose expected lookup time is $\Theta(1 + H)$.

Weaknesses of Static Optimality

- Statically optimal BSTs are fantastic if the lookups are sampled randomly from a fixed distribution.
- However, what if you don't know anything about the particular access pattern you're going to have?
- ***Question 1:*** Is it possible to build a BST with $O(1 + H)$ expected lookup time if the probability distribution isn't known in advance?

Weaknesses of Static Optimality

- Just knowing the access probabilities doesn't guarantee that you can build a good BST.
- ***Example:*** Suppose you're working at Yelp and want to support restaurant searches worldwide.
- There's some baseline probability distribution about which places get more lookups (New York City) than others (Barrow, AK).
- However, there's also a time-sensitivity: some areas will be "hotter" for searches than others based on wherever people are looking for lunches or dinners.
- ***Question 2:*** Can we build a BST that adapts to access patterns beyond just the global access probabilities?

Challenge: Can we build a BST that meets the static optimality requirements, but is also sensitive to access patterns?

And can we do it without advance knowledge of the access pattern?

The Intuition

- If we don't know the access probabilities in advance, we can't build a fixed BST and then “hope” it works correctly.
- Instead, we'll have to restructure the BST as operations are performed.
- For now, let's focus on lookups; we'll handle insertions and deletions later on.

Refresher: Tree Rotations

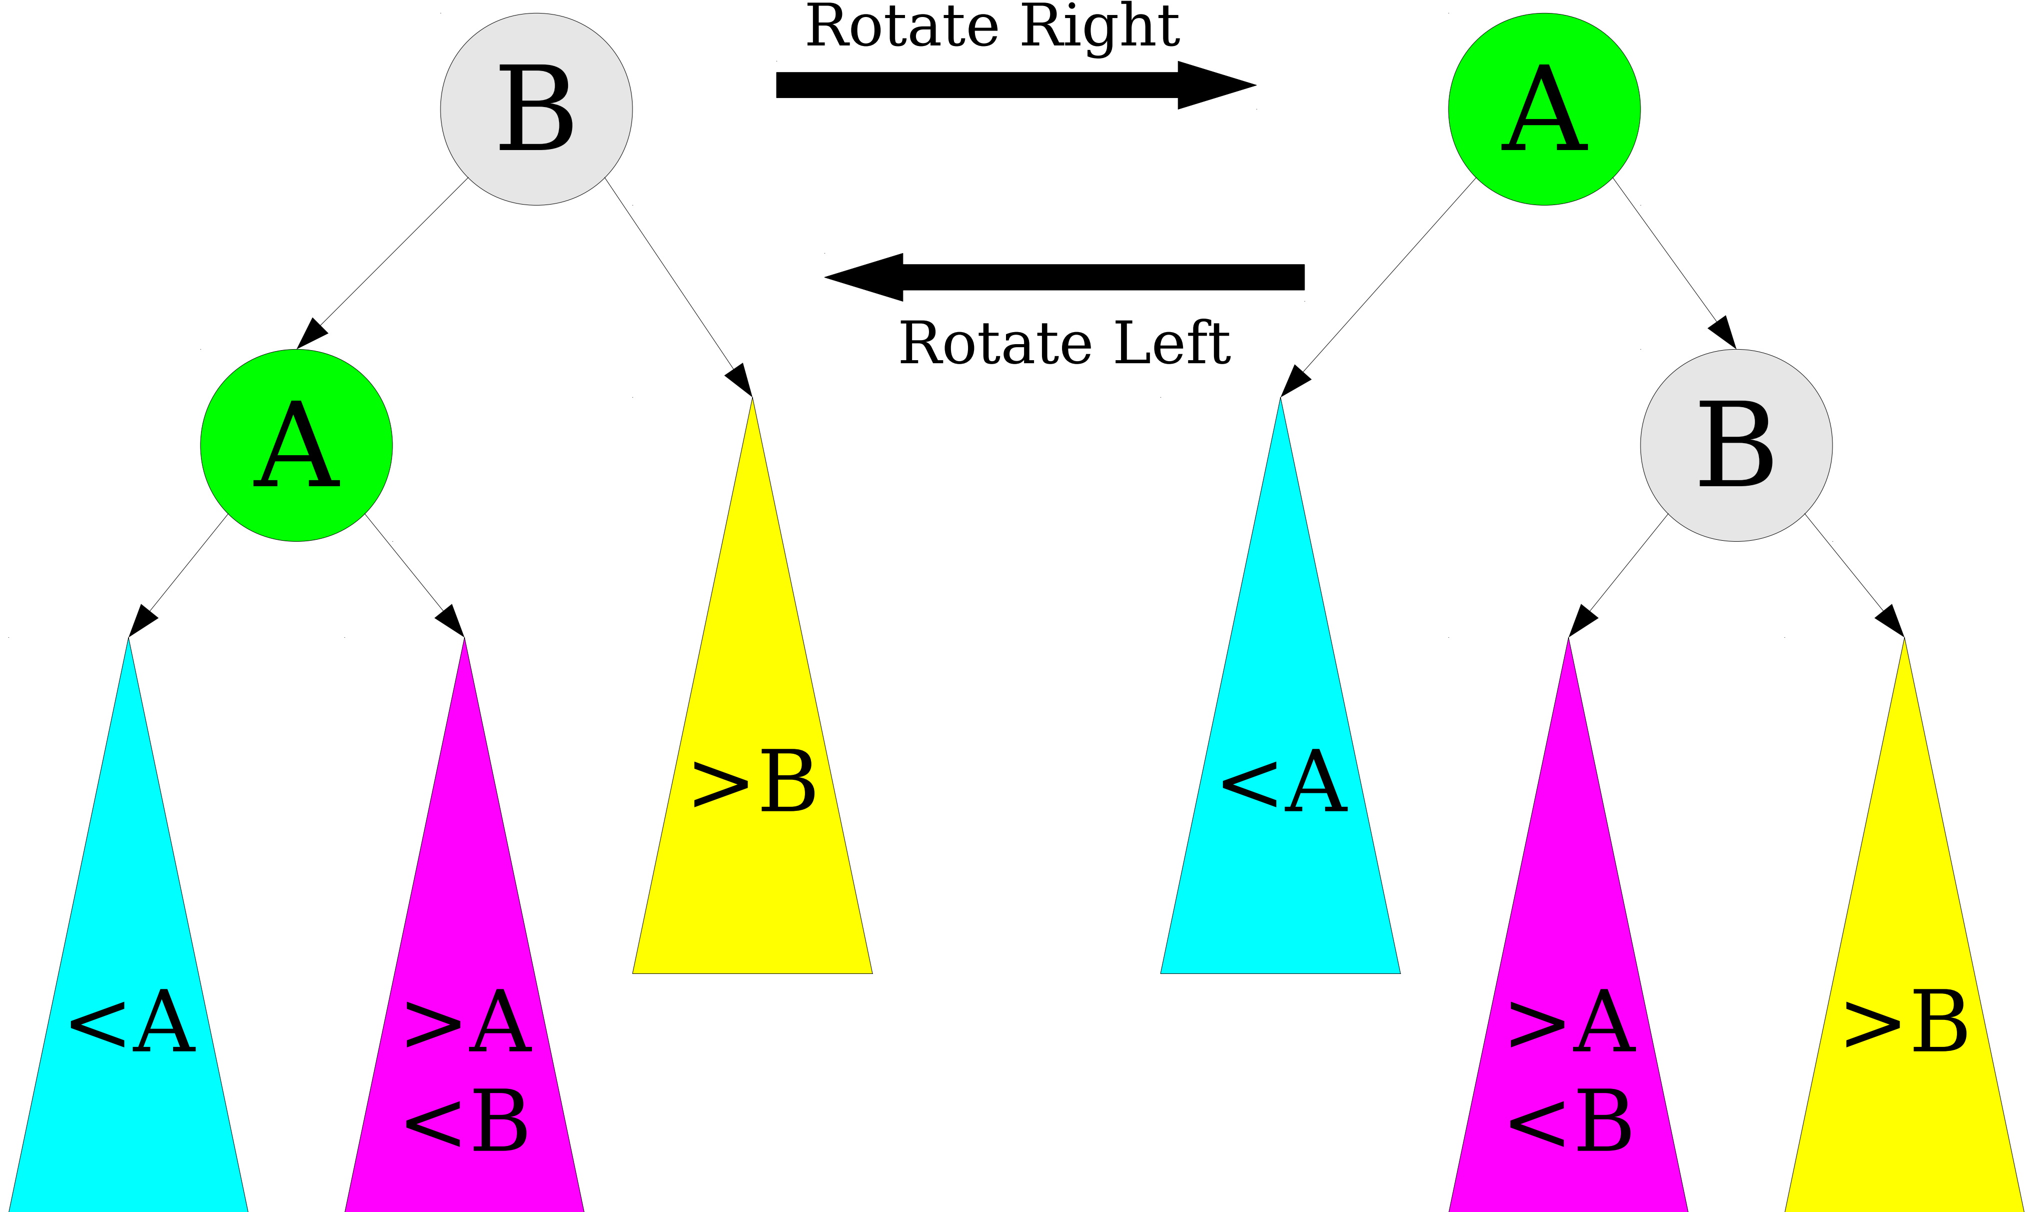

An Initial Approach

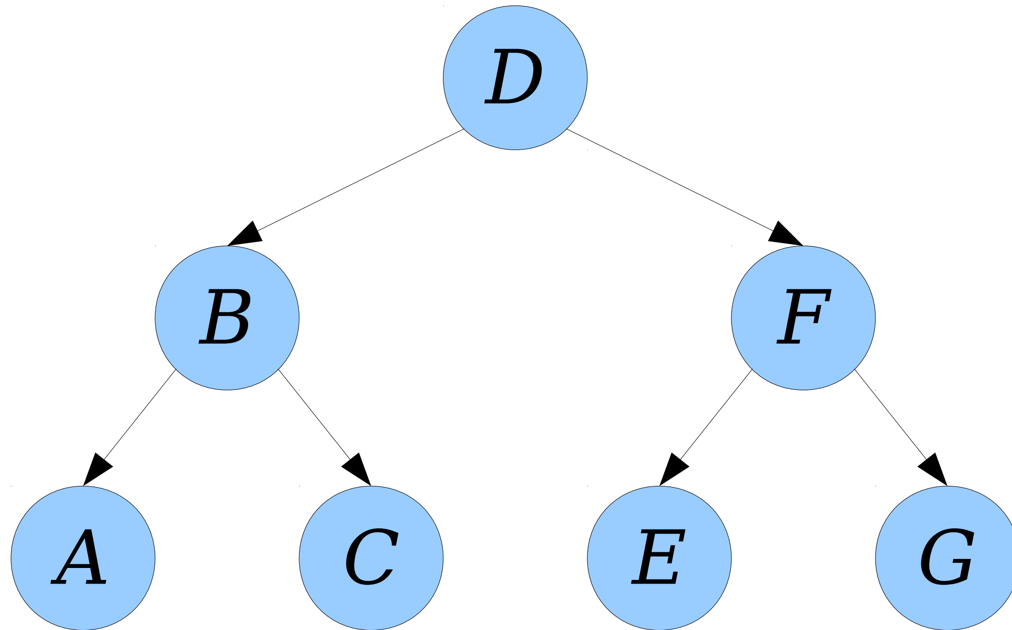

An Initial Approach

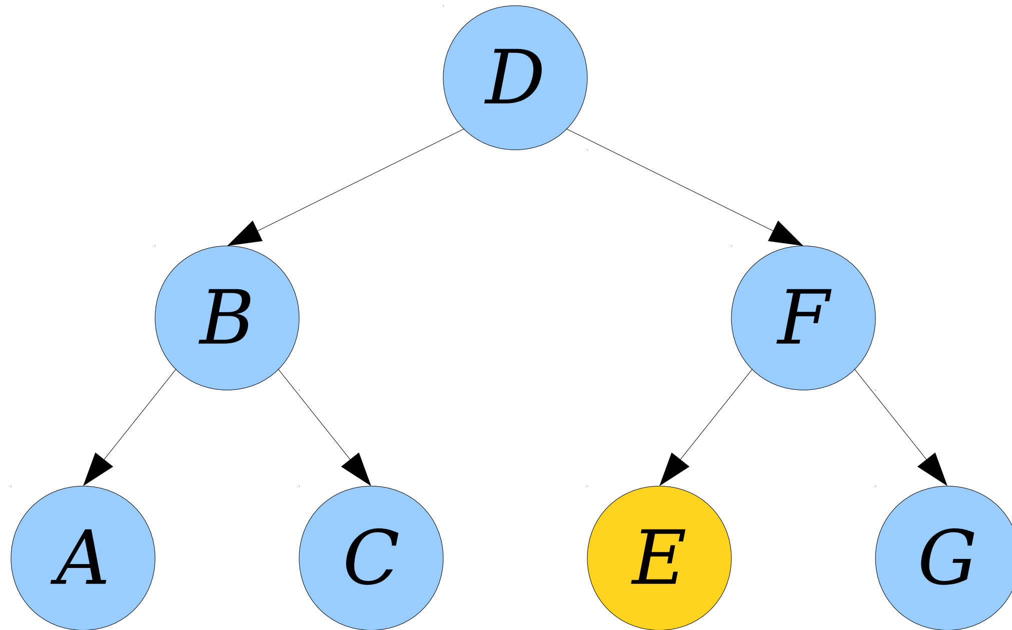

An Initial Approach

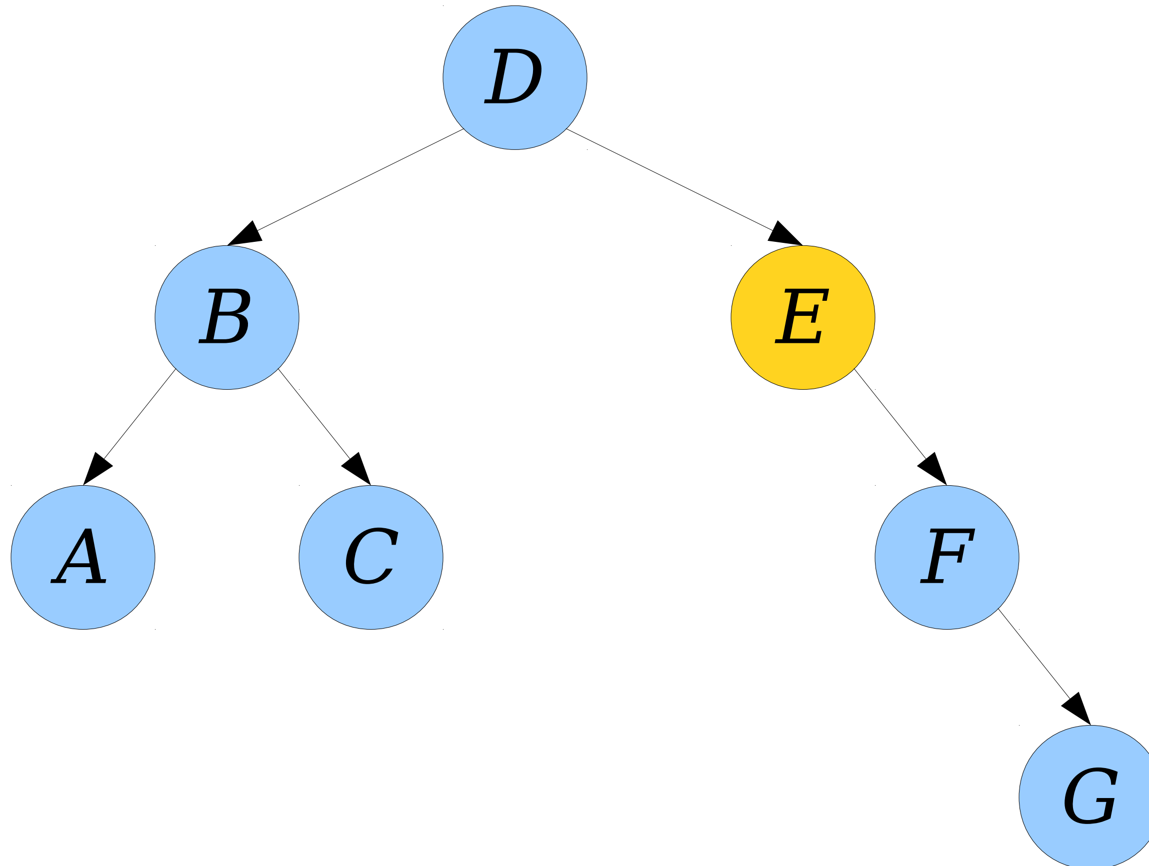

An Initial Approach

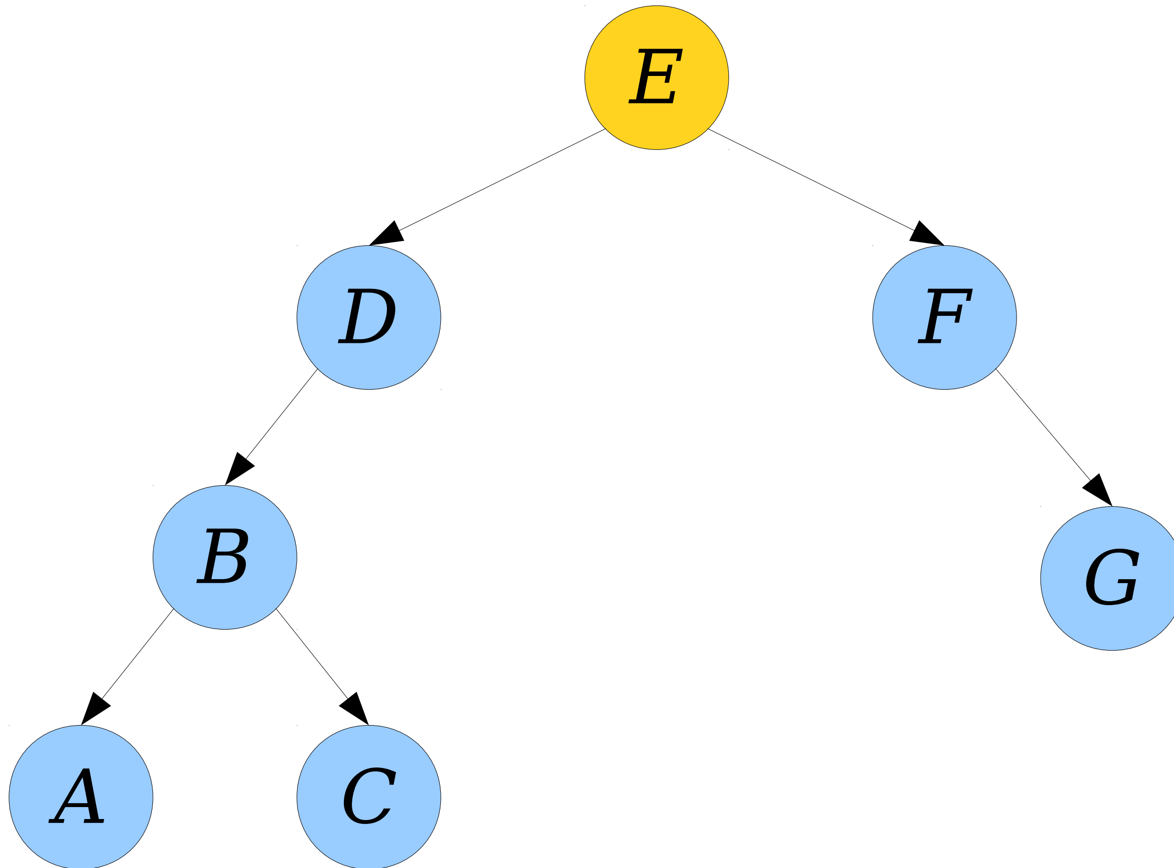

An Initial Approach

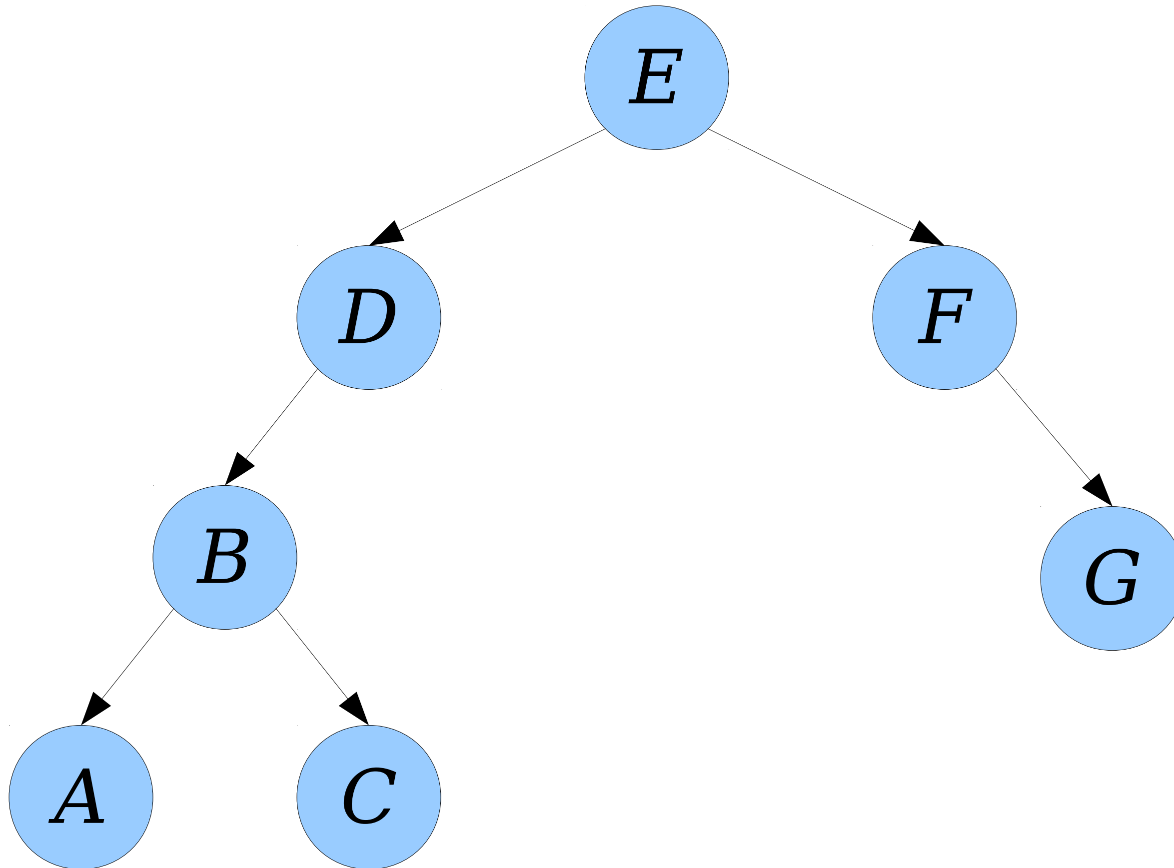

An Initial Approach

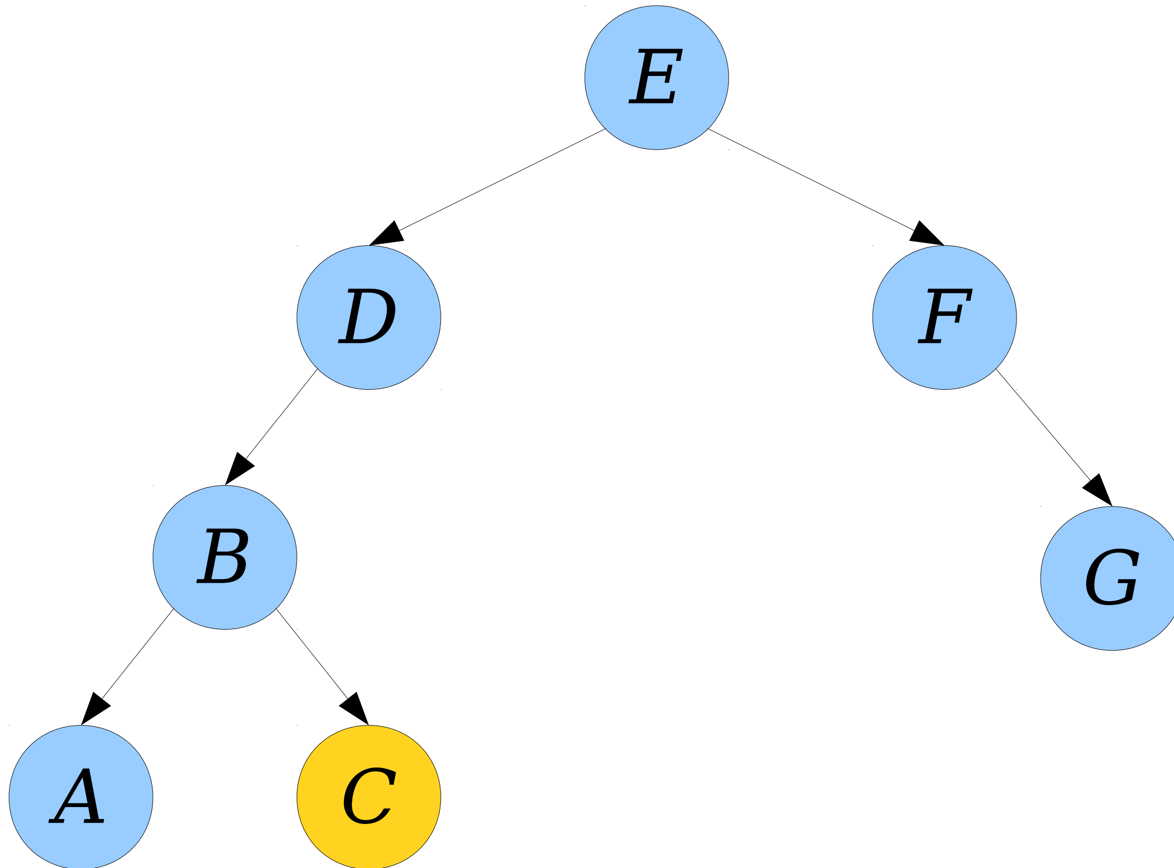

An Initial Approach

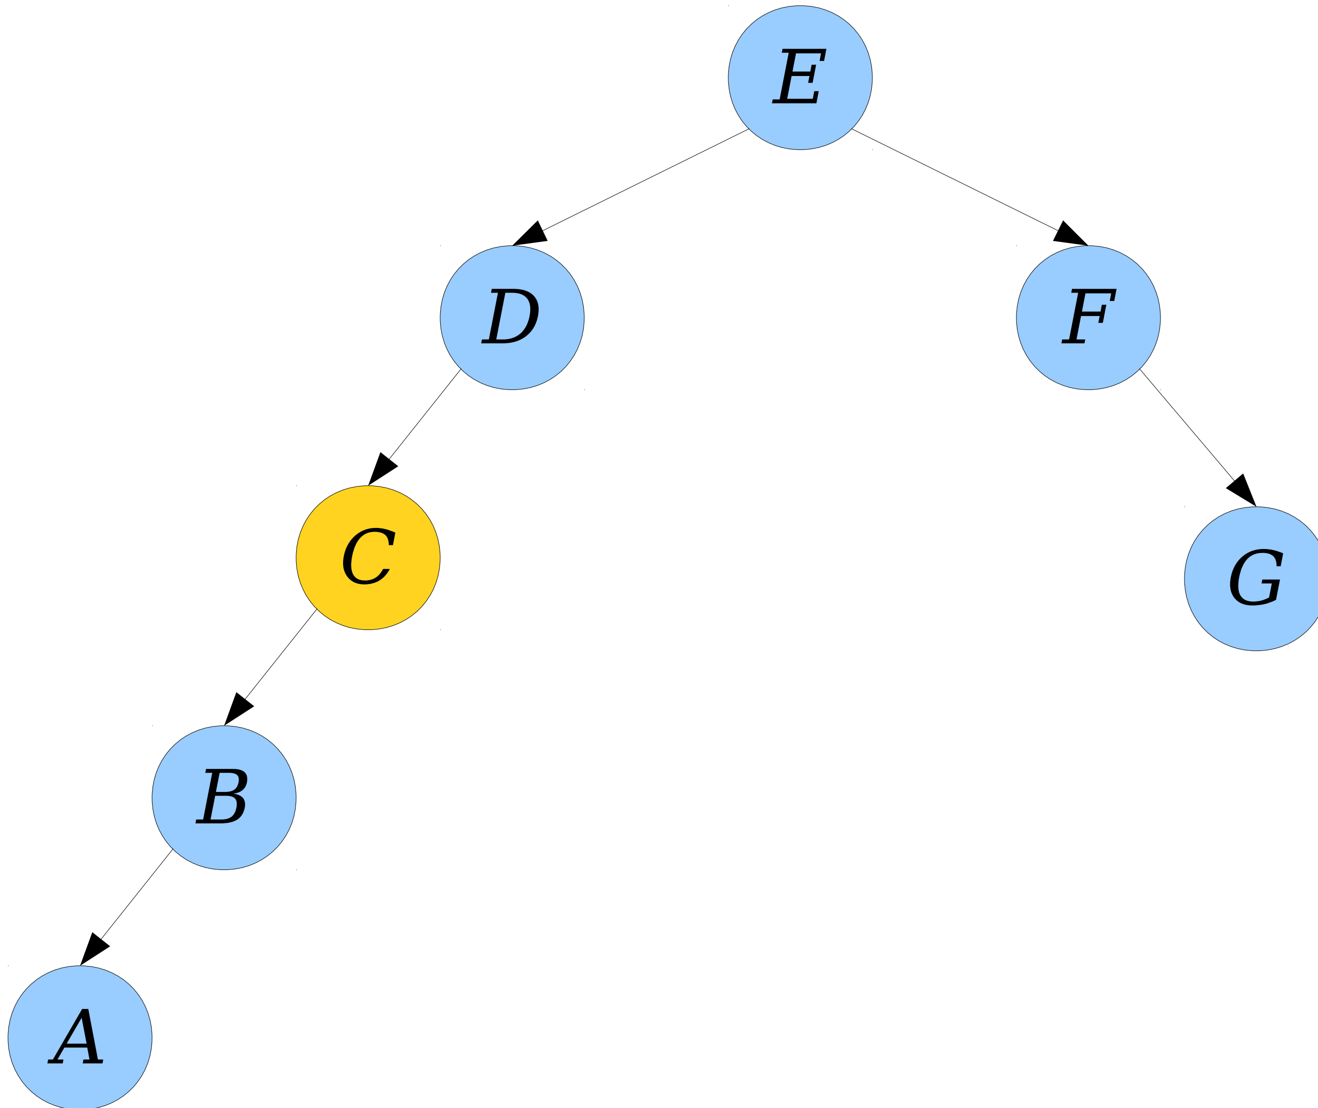

An Initial Approach

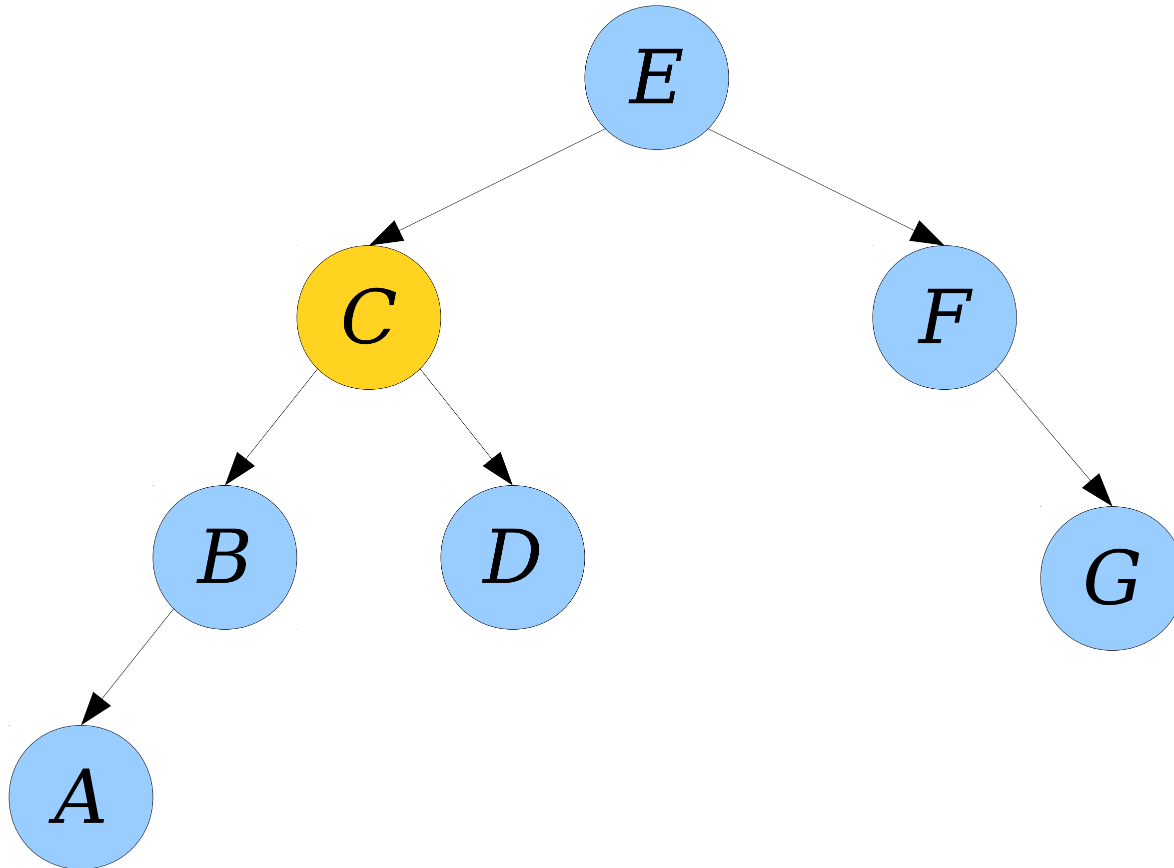

An Initial Approach

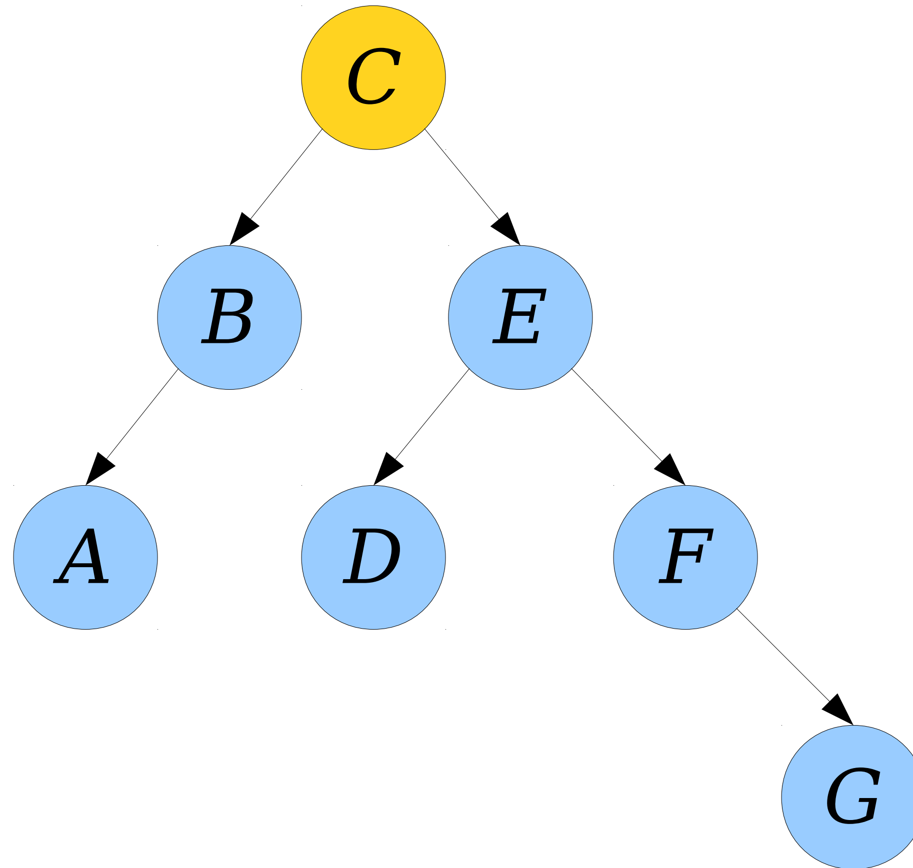

An Initial Approach

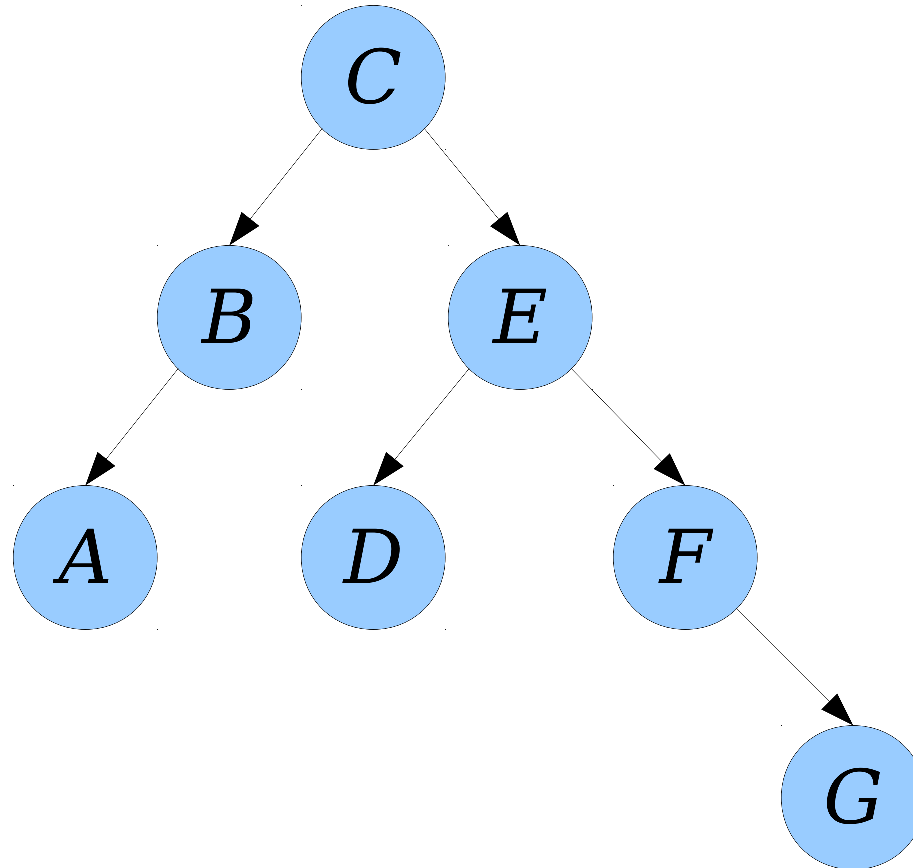

An Initial Approach

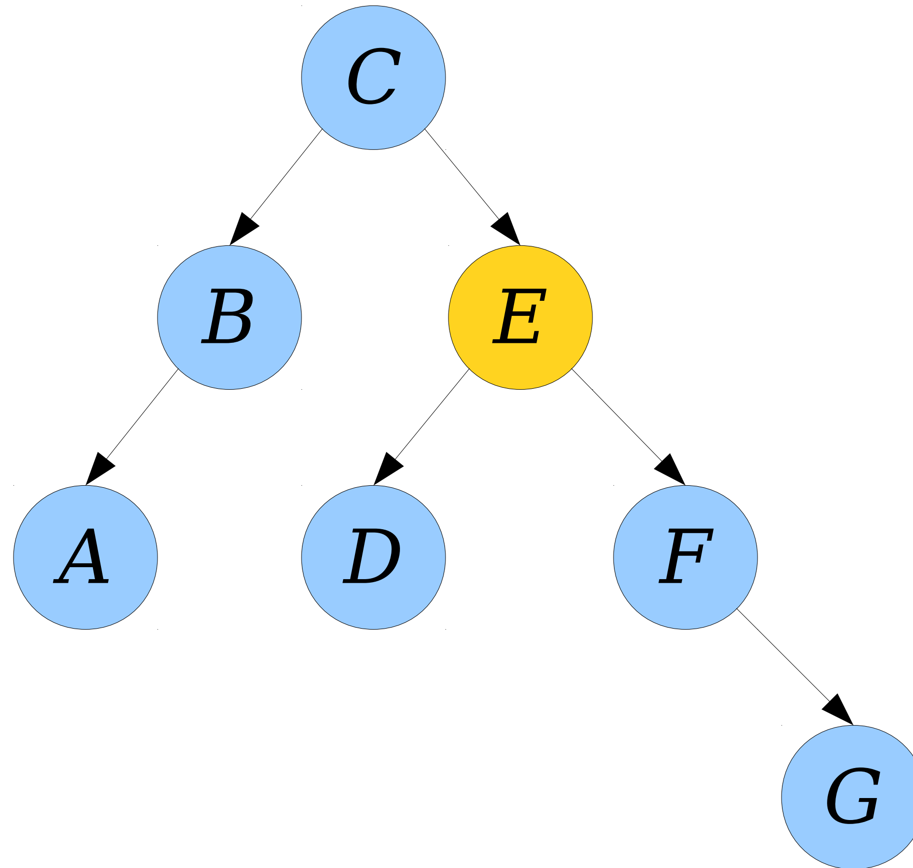

An Initial Approach

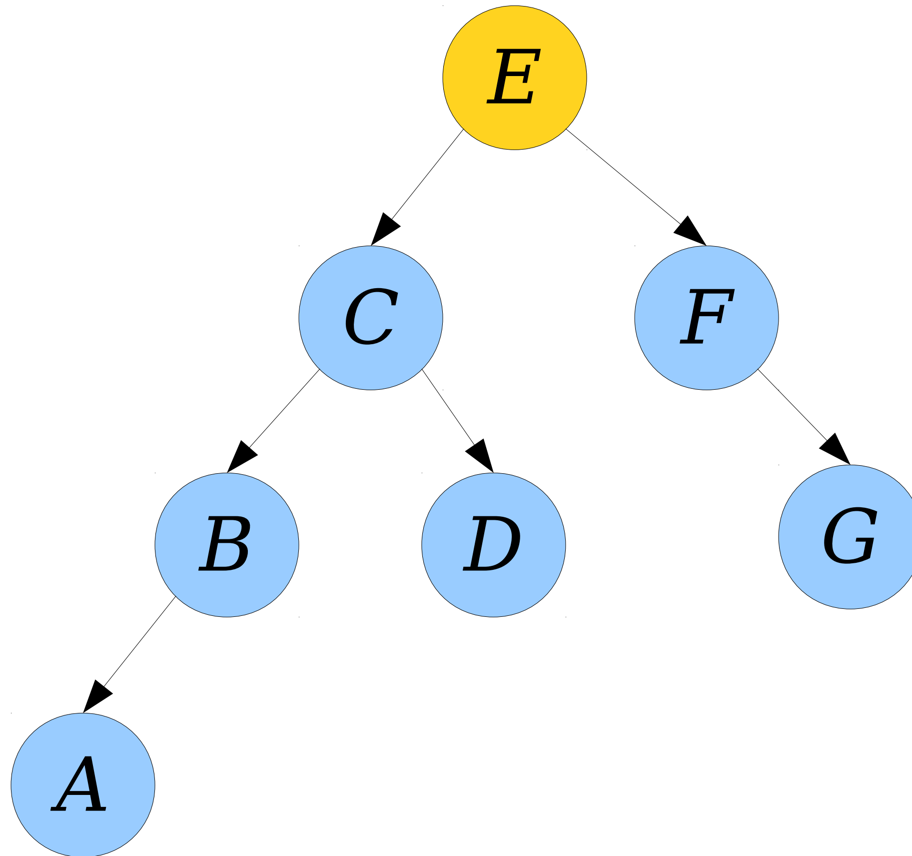

An Initial Idea

- Begin with an arbitrary BST.
- After looking up an element, repeatedly rotate that element with its parent until it becomes the root.
- ***Intuition:***
    - Recently-accessed elements will be up near the root of the tree, lowering access time.
    - Unused elements stay low in the tree.

The Problem

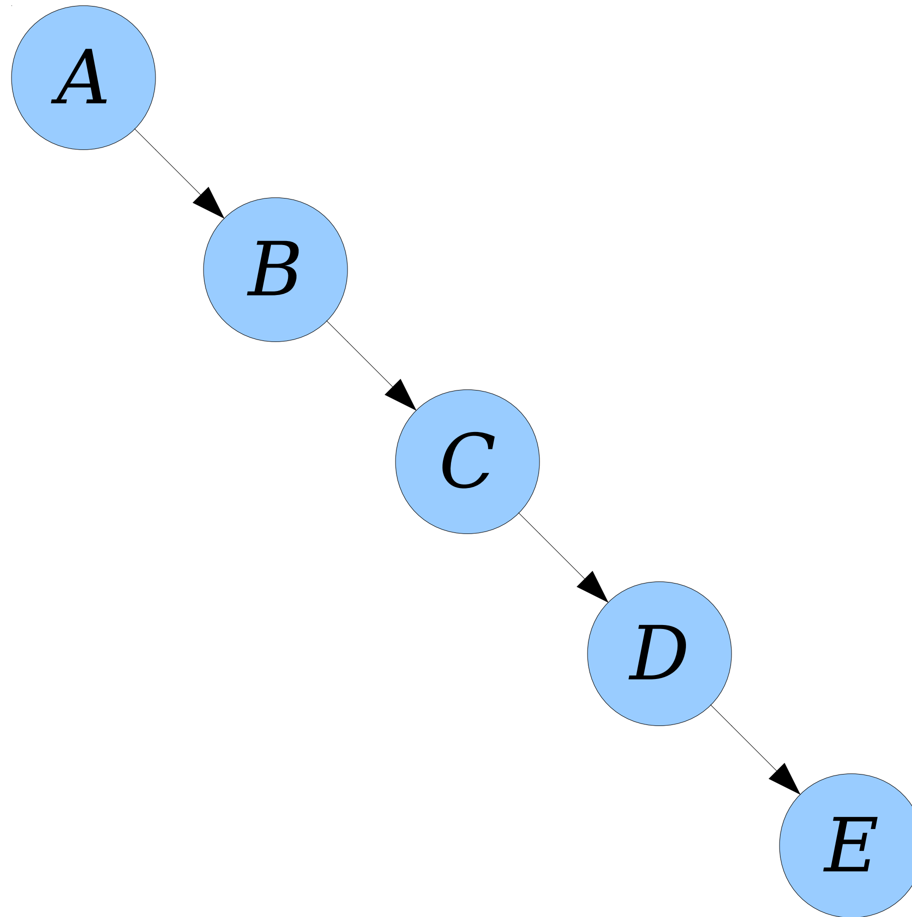

The Problem

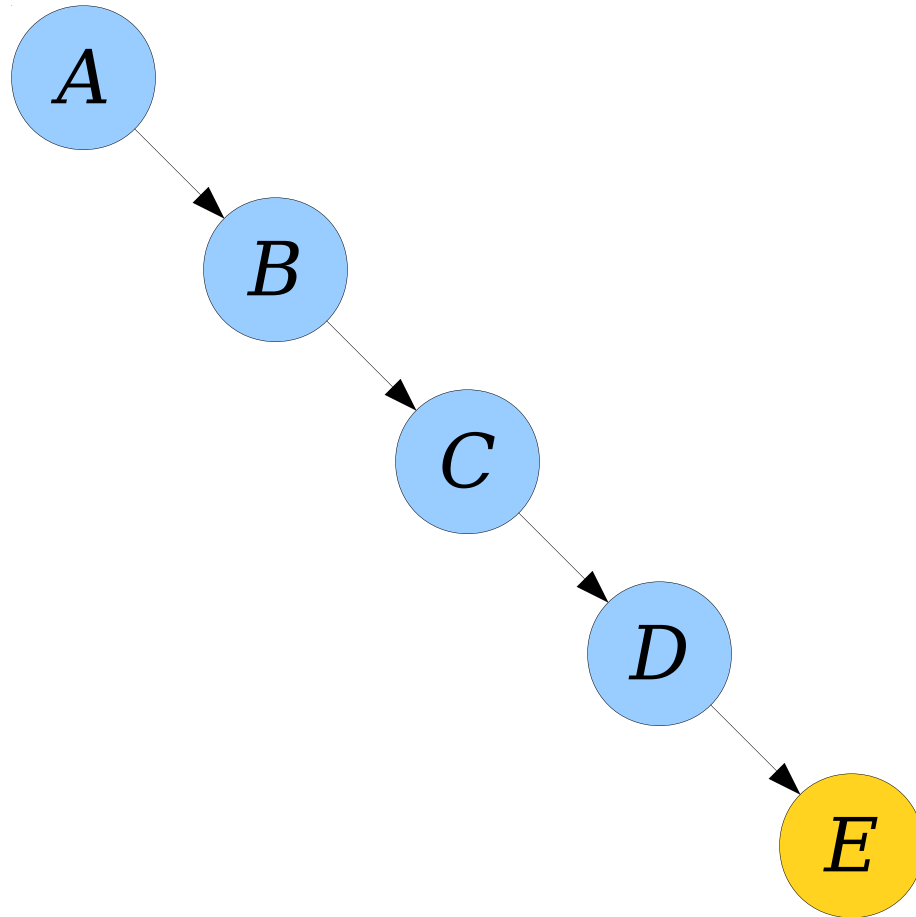

The Problem

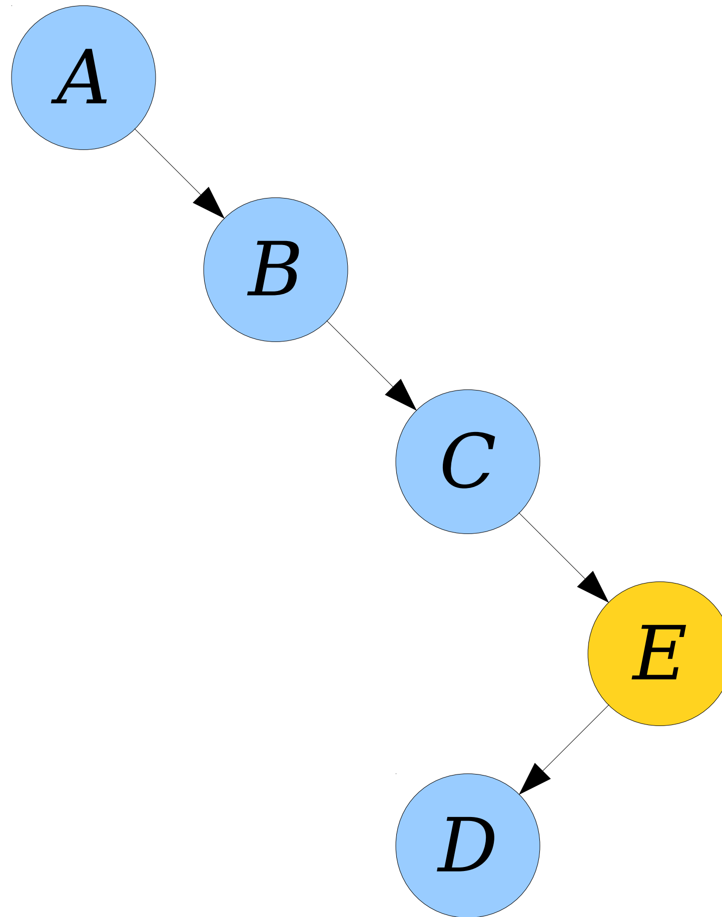

The Problem

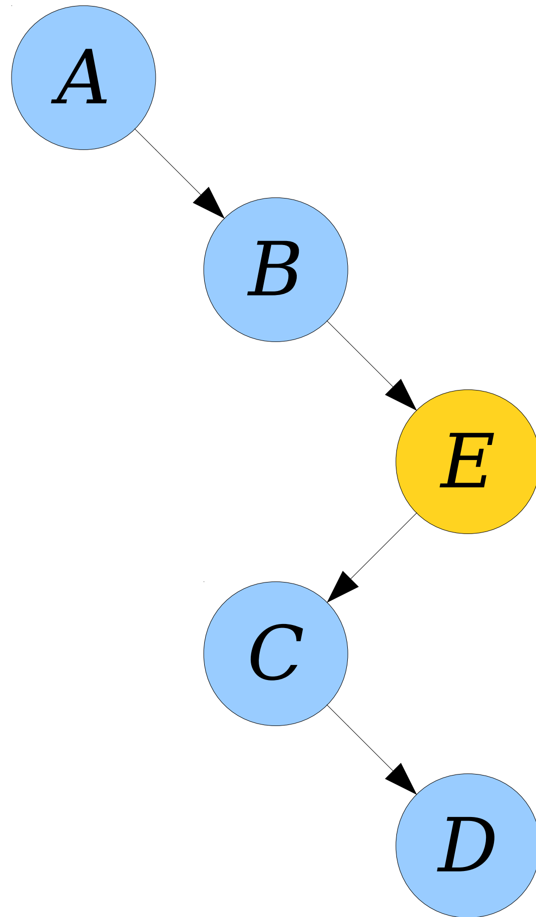

The Problem

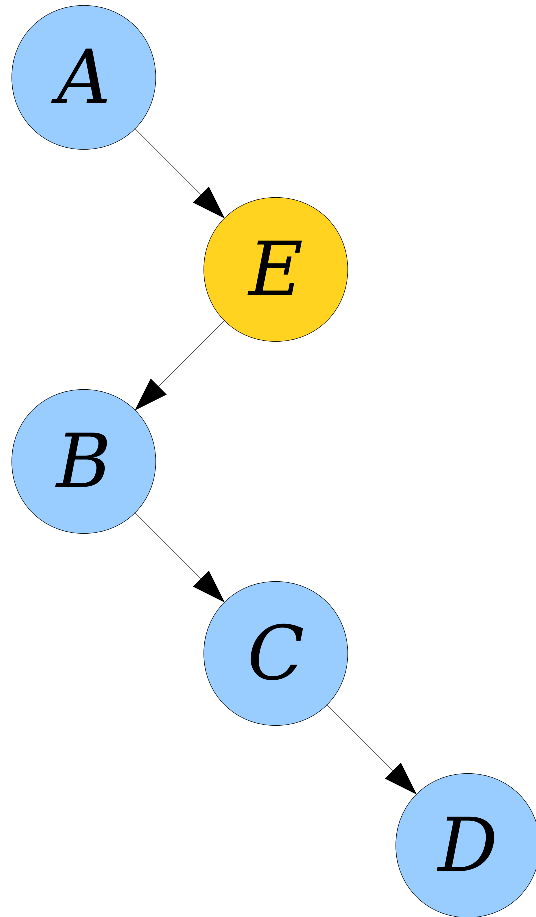

The Problem

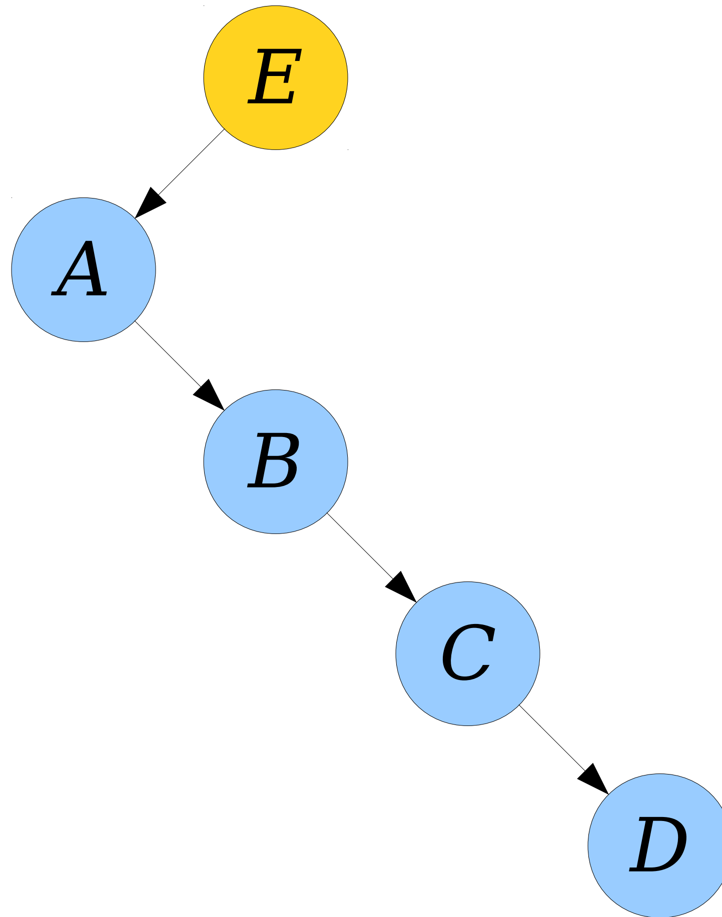

The Problem

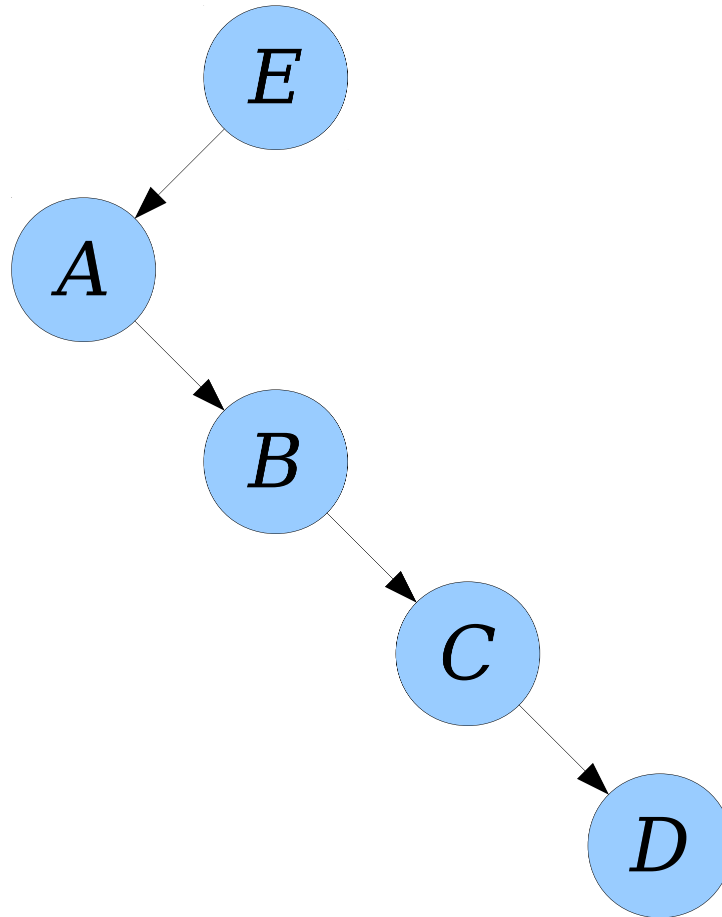

The Problem

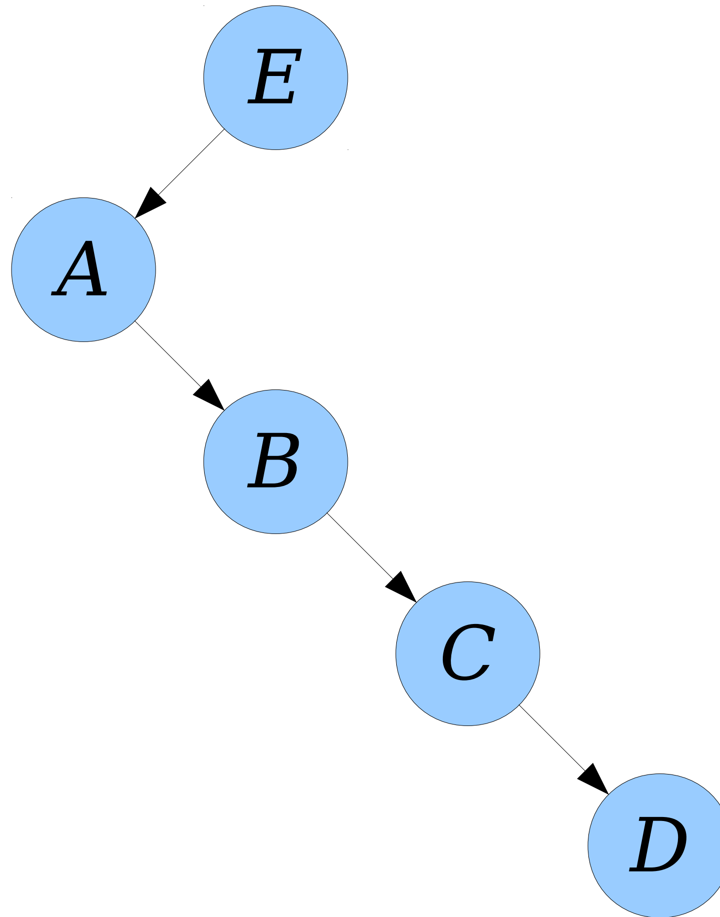

Rotations Needed: **5**

The Problem

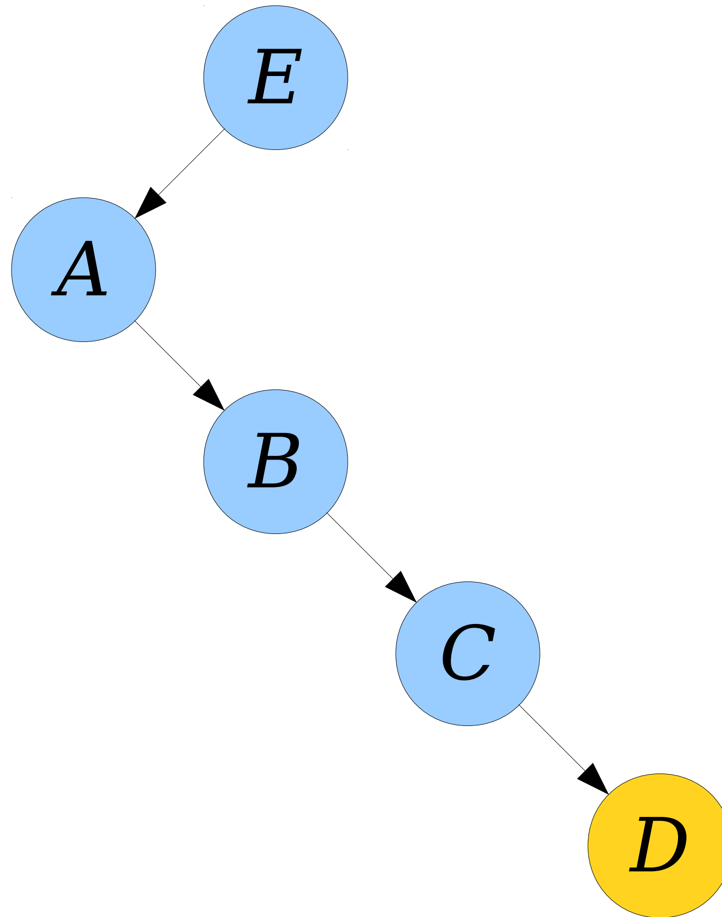

The Problem

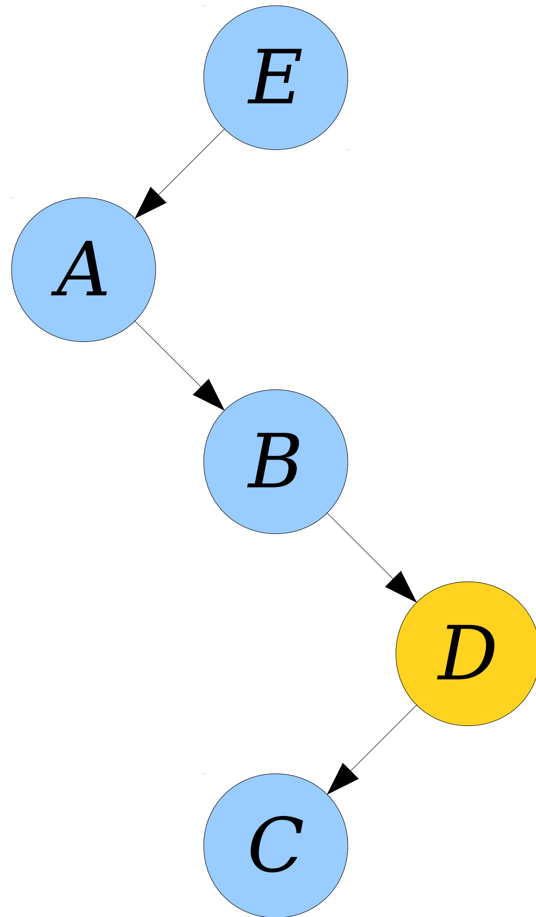

The Problem

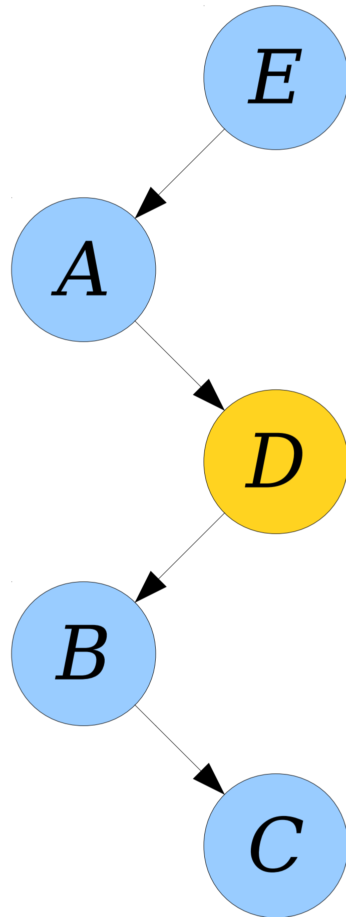

The Problem

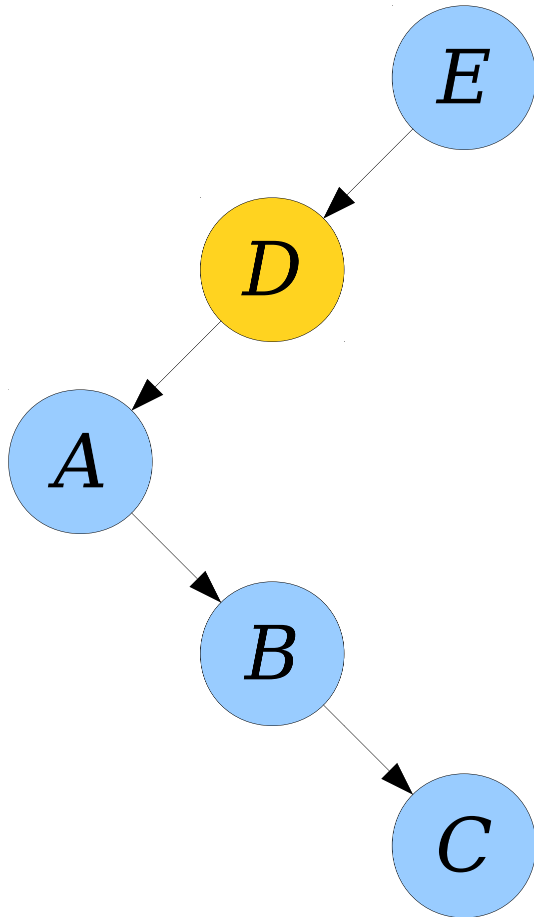

The Problem

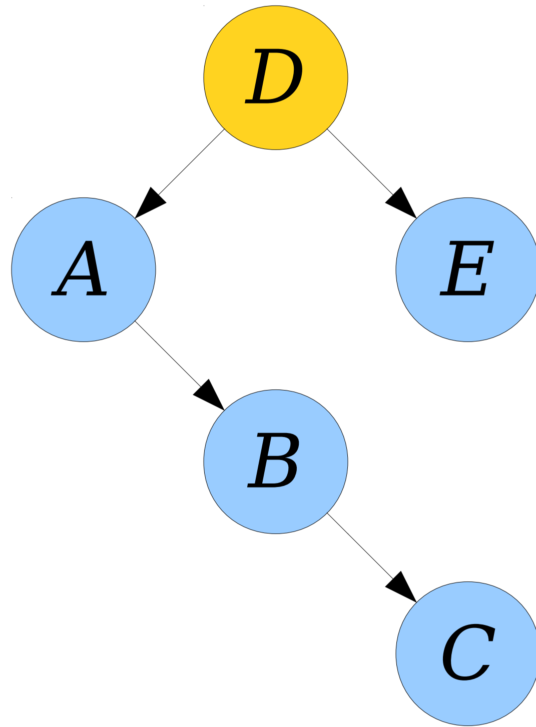

The Problem

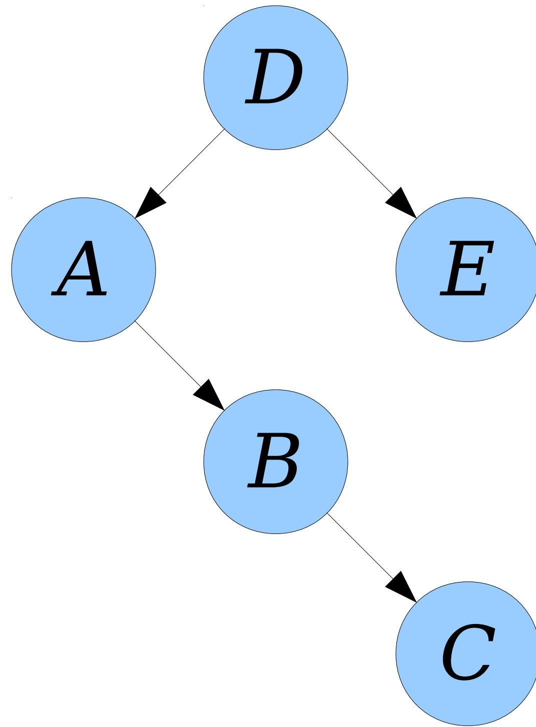

The Problem

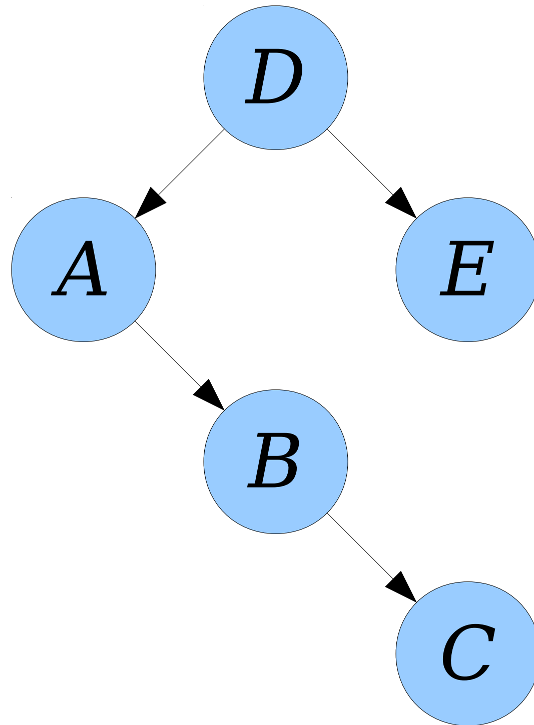

Rotations Needed: **4**

The Problem

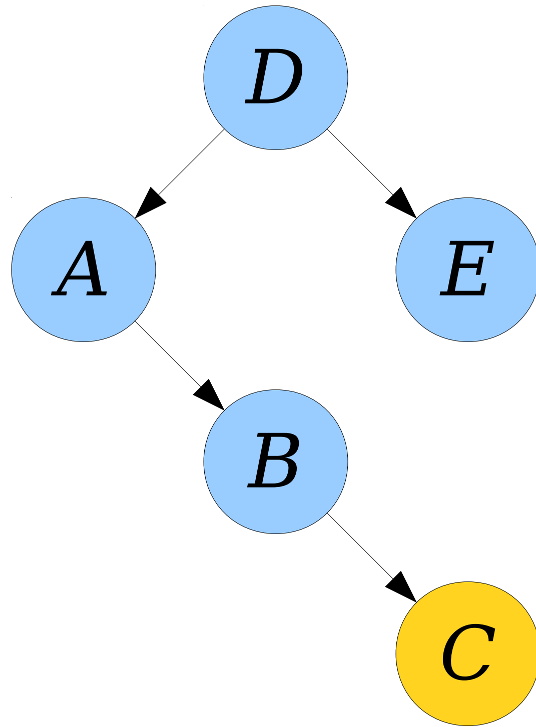

The Problem

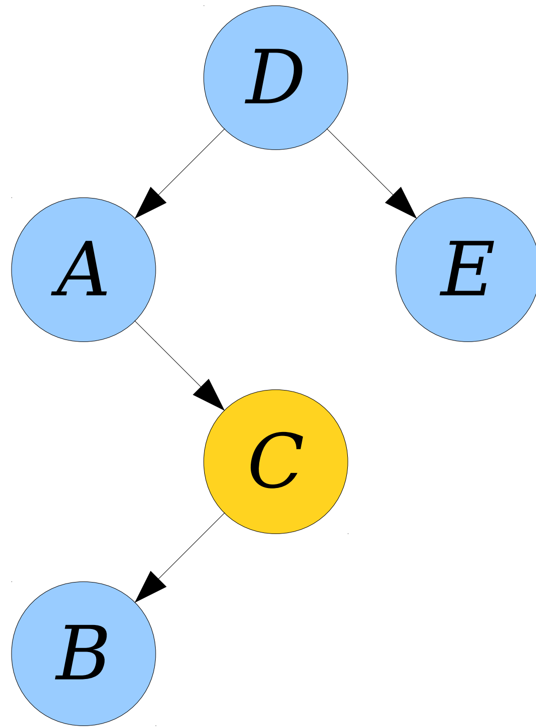

The Problem

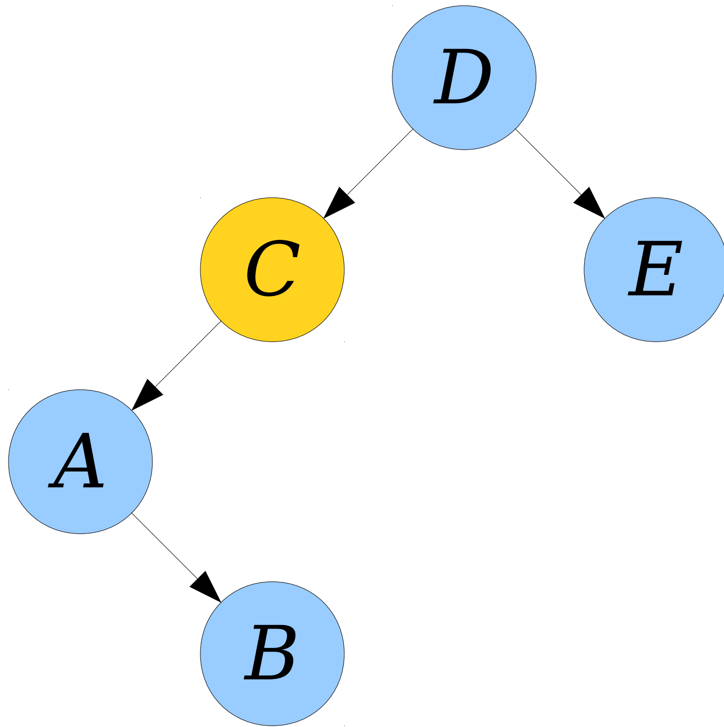

The Problem

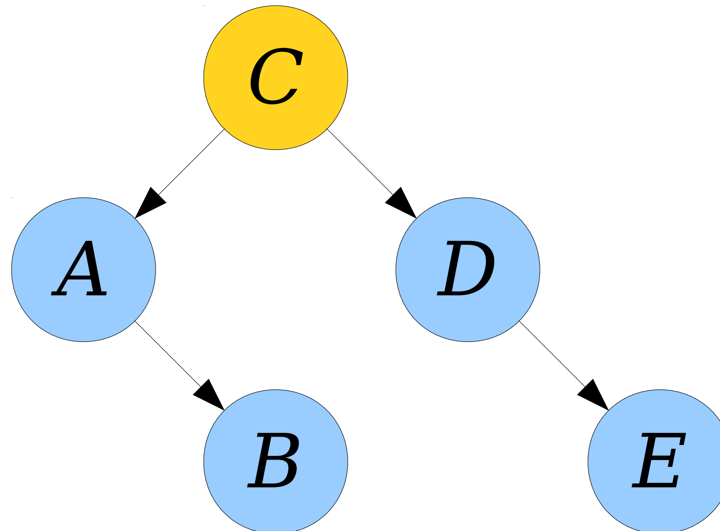

The Problem

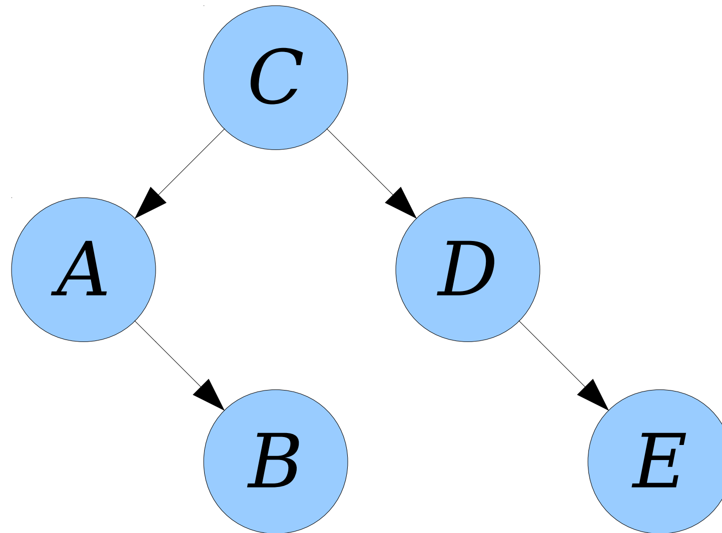

The Problem

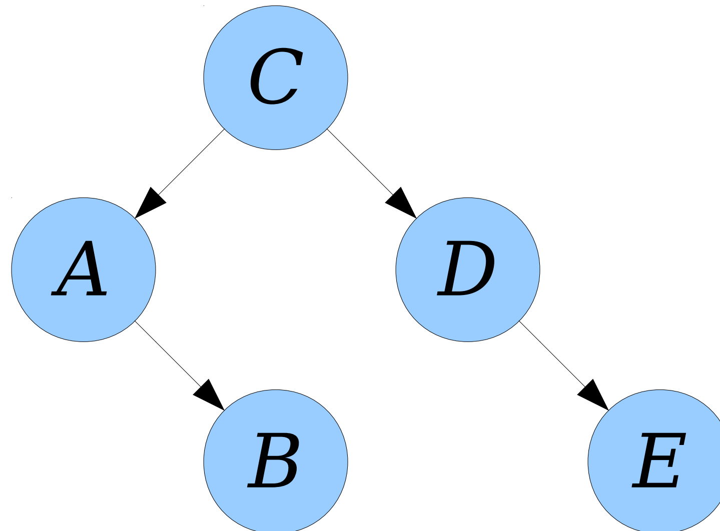

Rotations Needed: **3**

The Problem

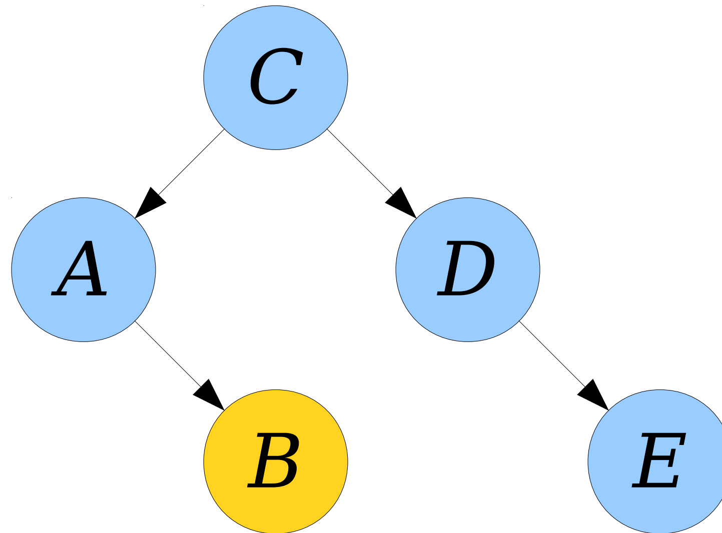

The Problem

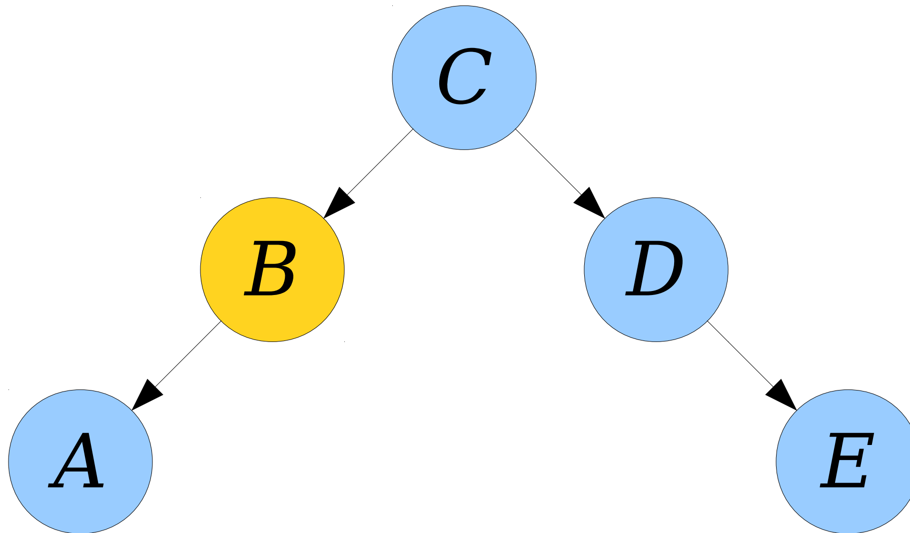

The Problem

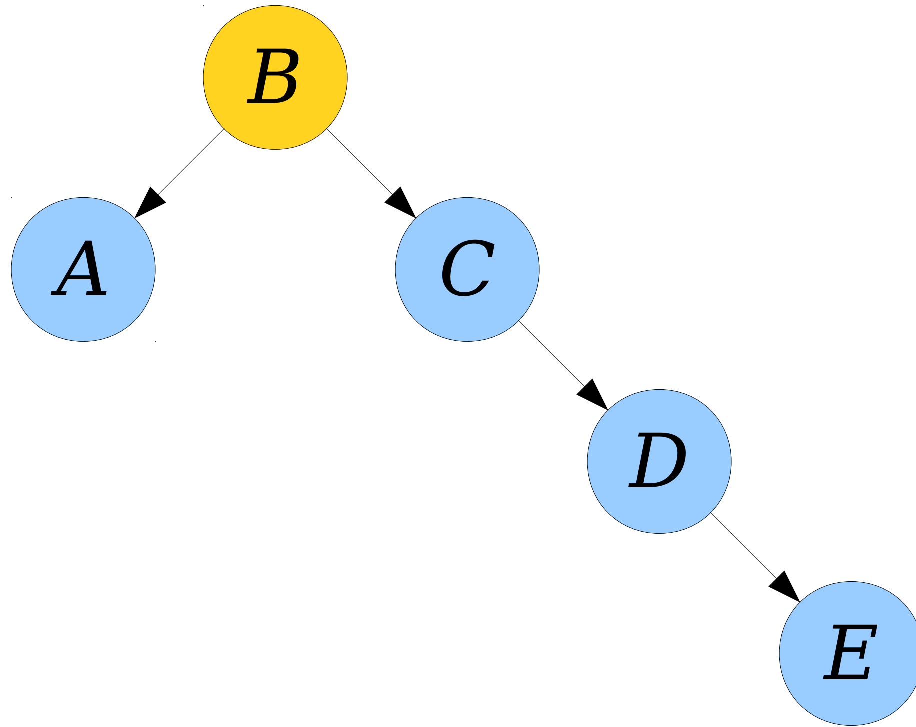

The Problem

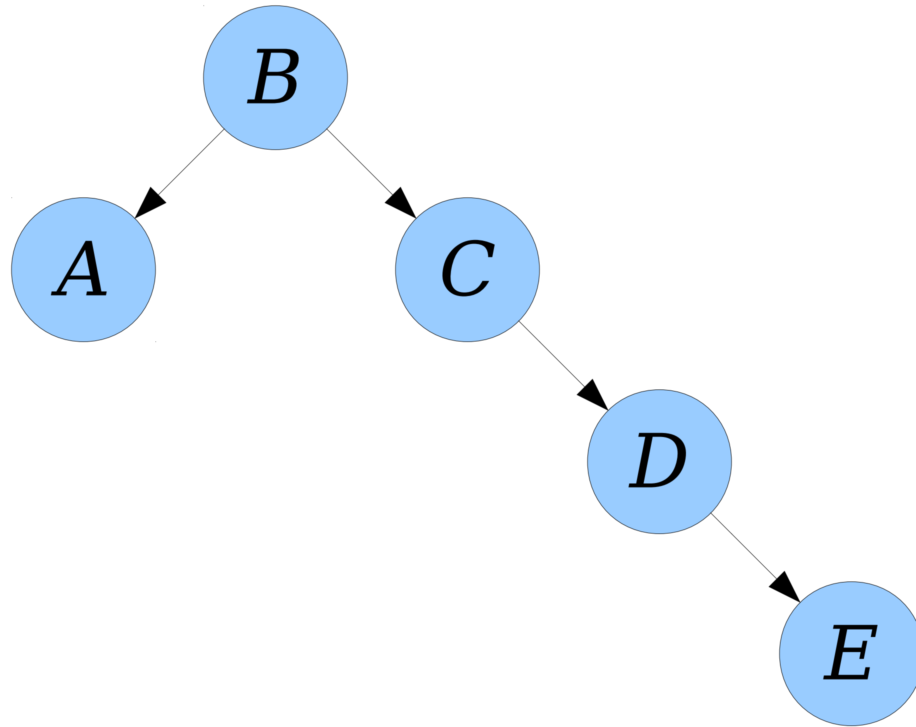

The Problem

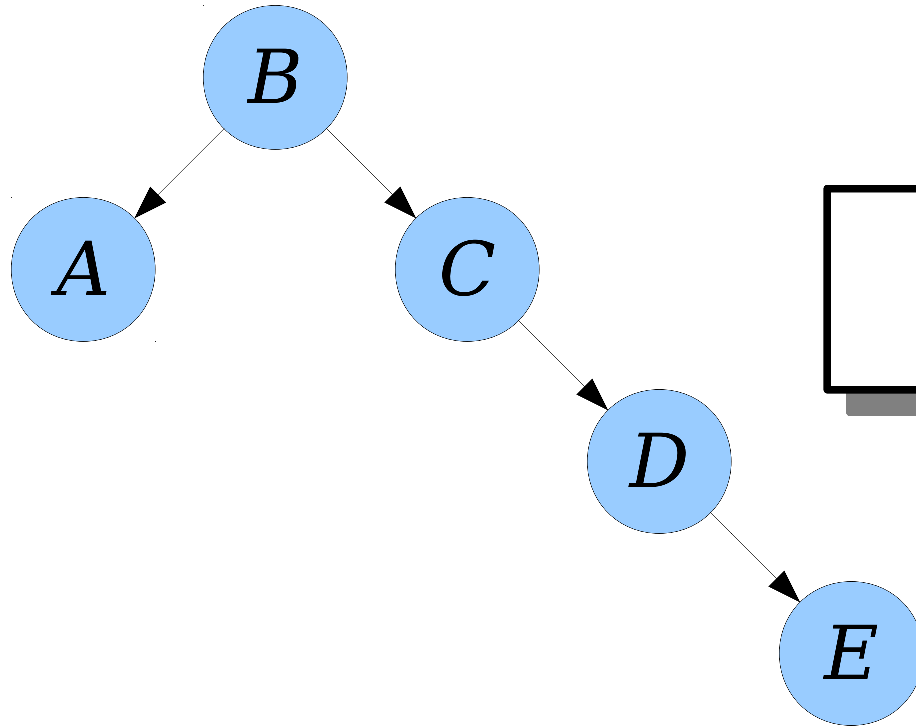

Rotations Needed: **2**

The Problem

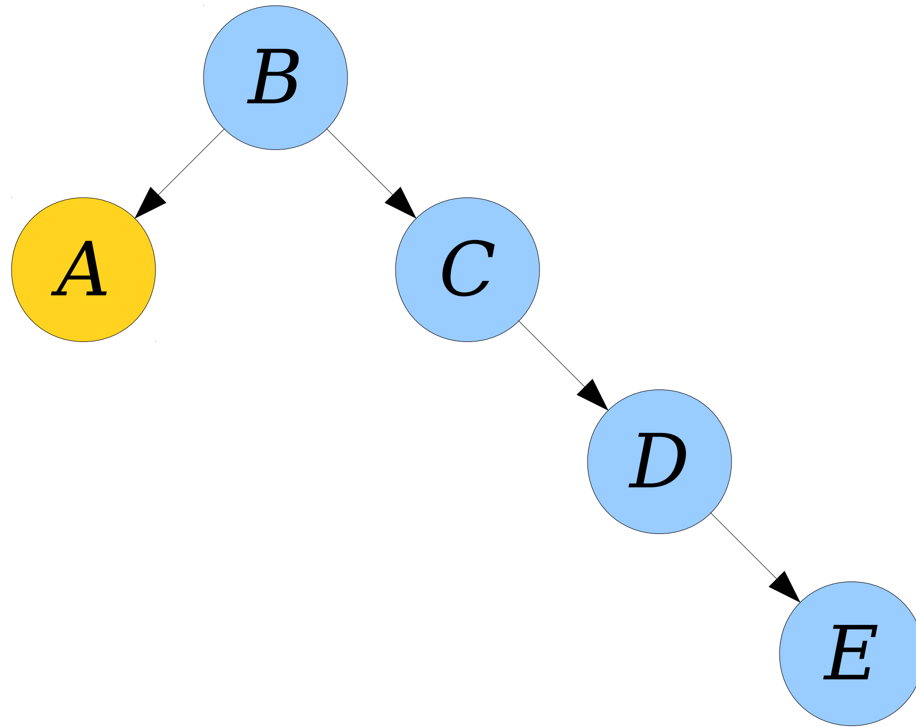

The Problem

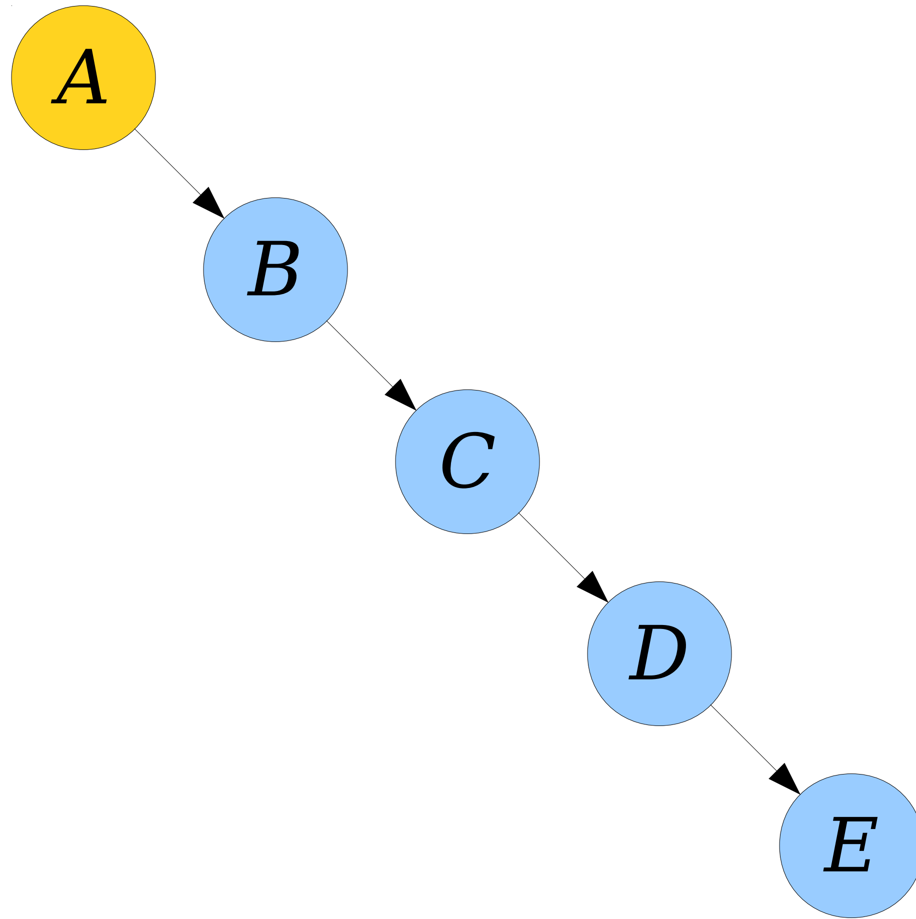

The Problem

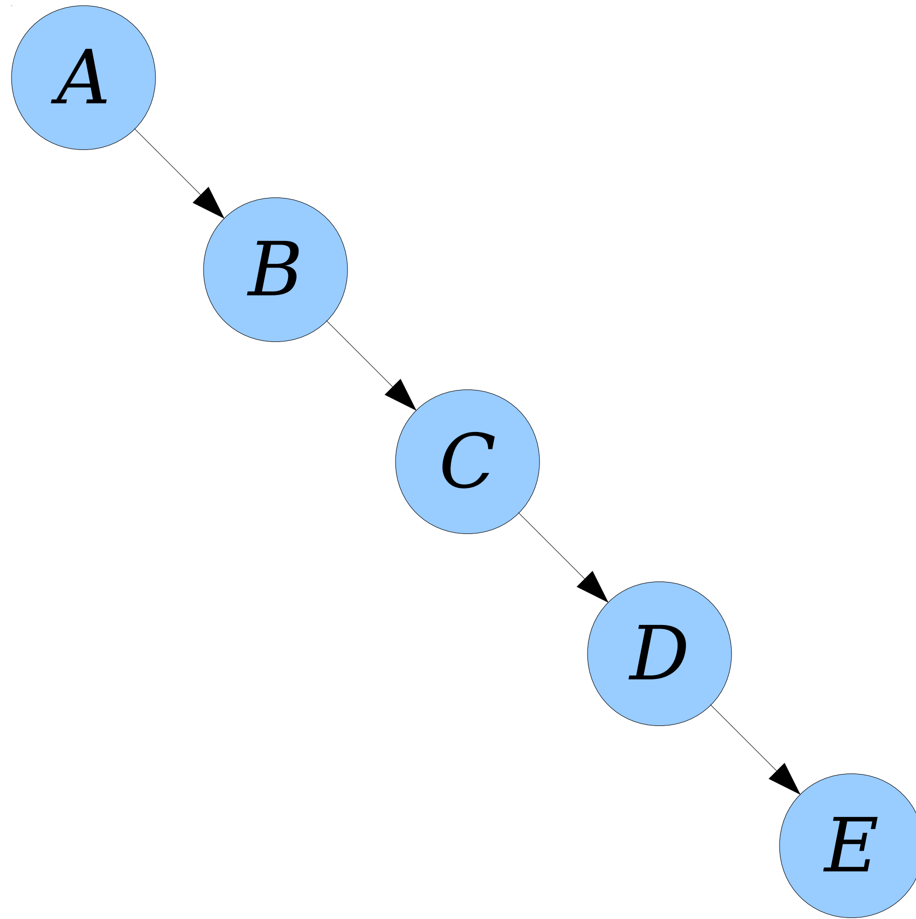

The Problem

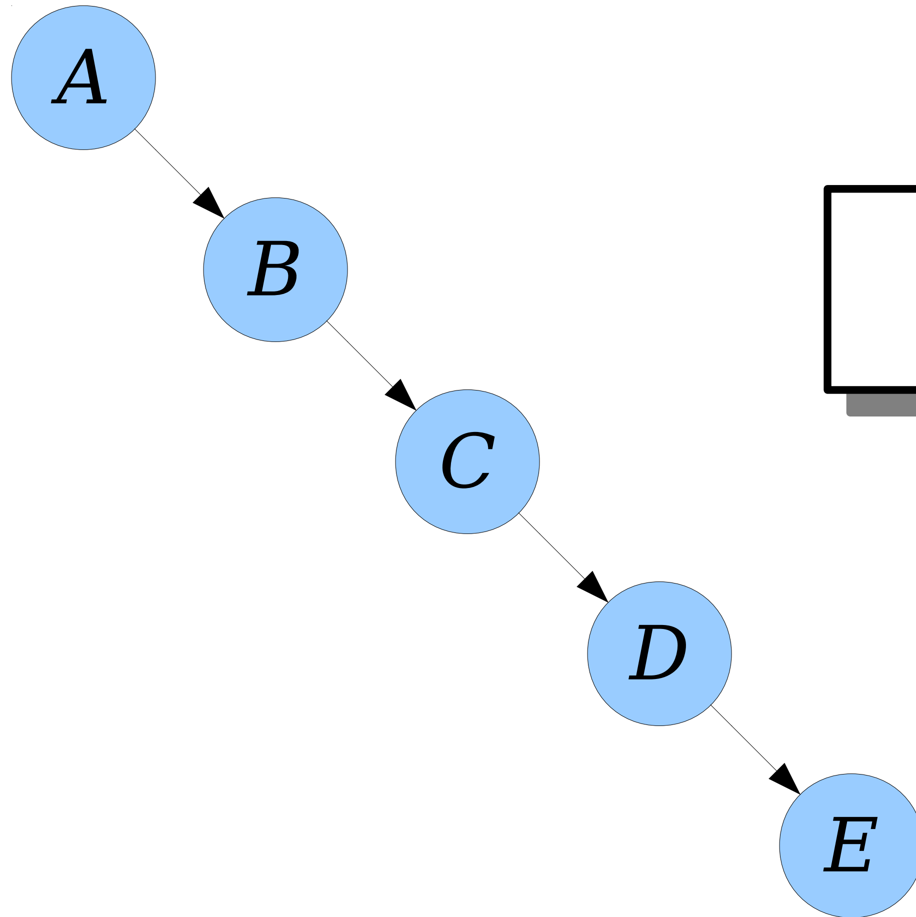

Rotations Needed: **1**

The Problem

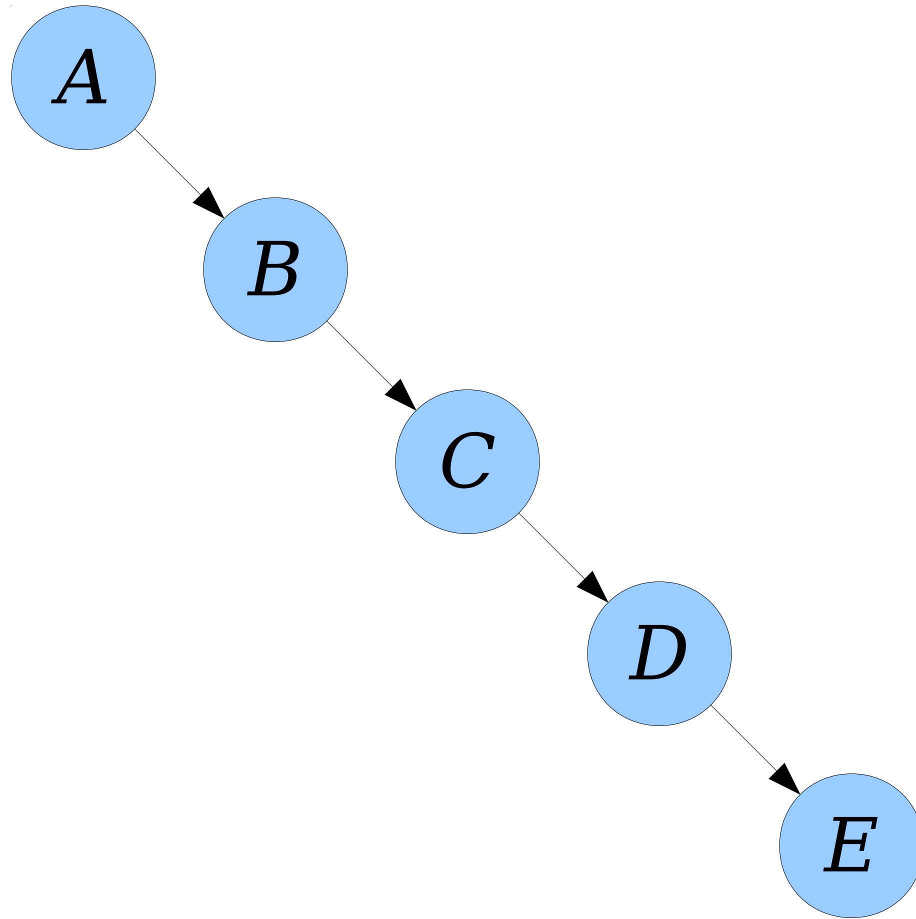

The Problem

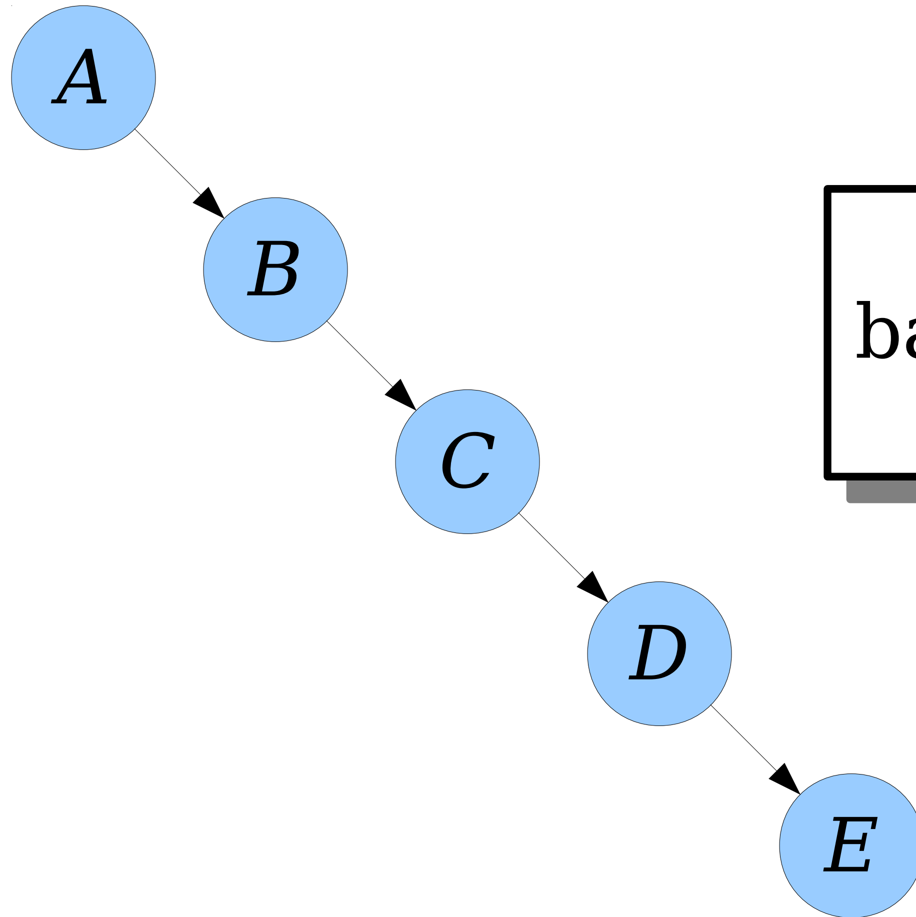

We're right back where we started!

The Problem

- The “rotate to root” method might result in n accesses taking time $\Theta(n^2)$.
- ***Why?***
- The cost of looking up a key x depends on the length of the access path to x.
- Rotating x up to the root doesn’t always improve that access path.
- Future lookups deep down near where x used to be will still be slow.

A More Balanced Approach

- In 1983, Daniel Sleator and Robert Tarjan invented an operation called ***splaying***.
- Splaying rotates an element to the root of the tree, but does so in a way that's more “fair” to other nodes in the tree.
- Each splay works by applying one of three templates to determine which rotations to apply.

Case 1: Zig-Zag

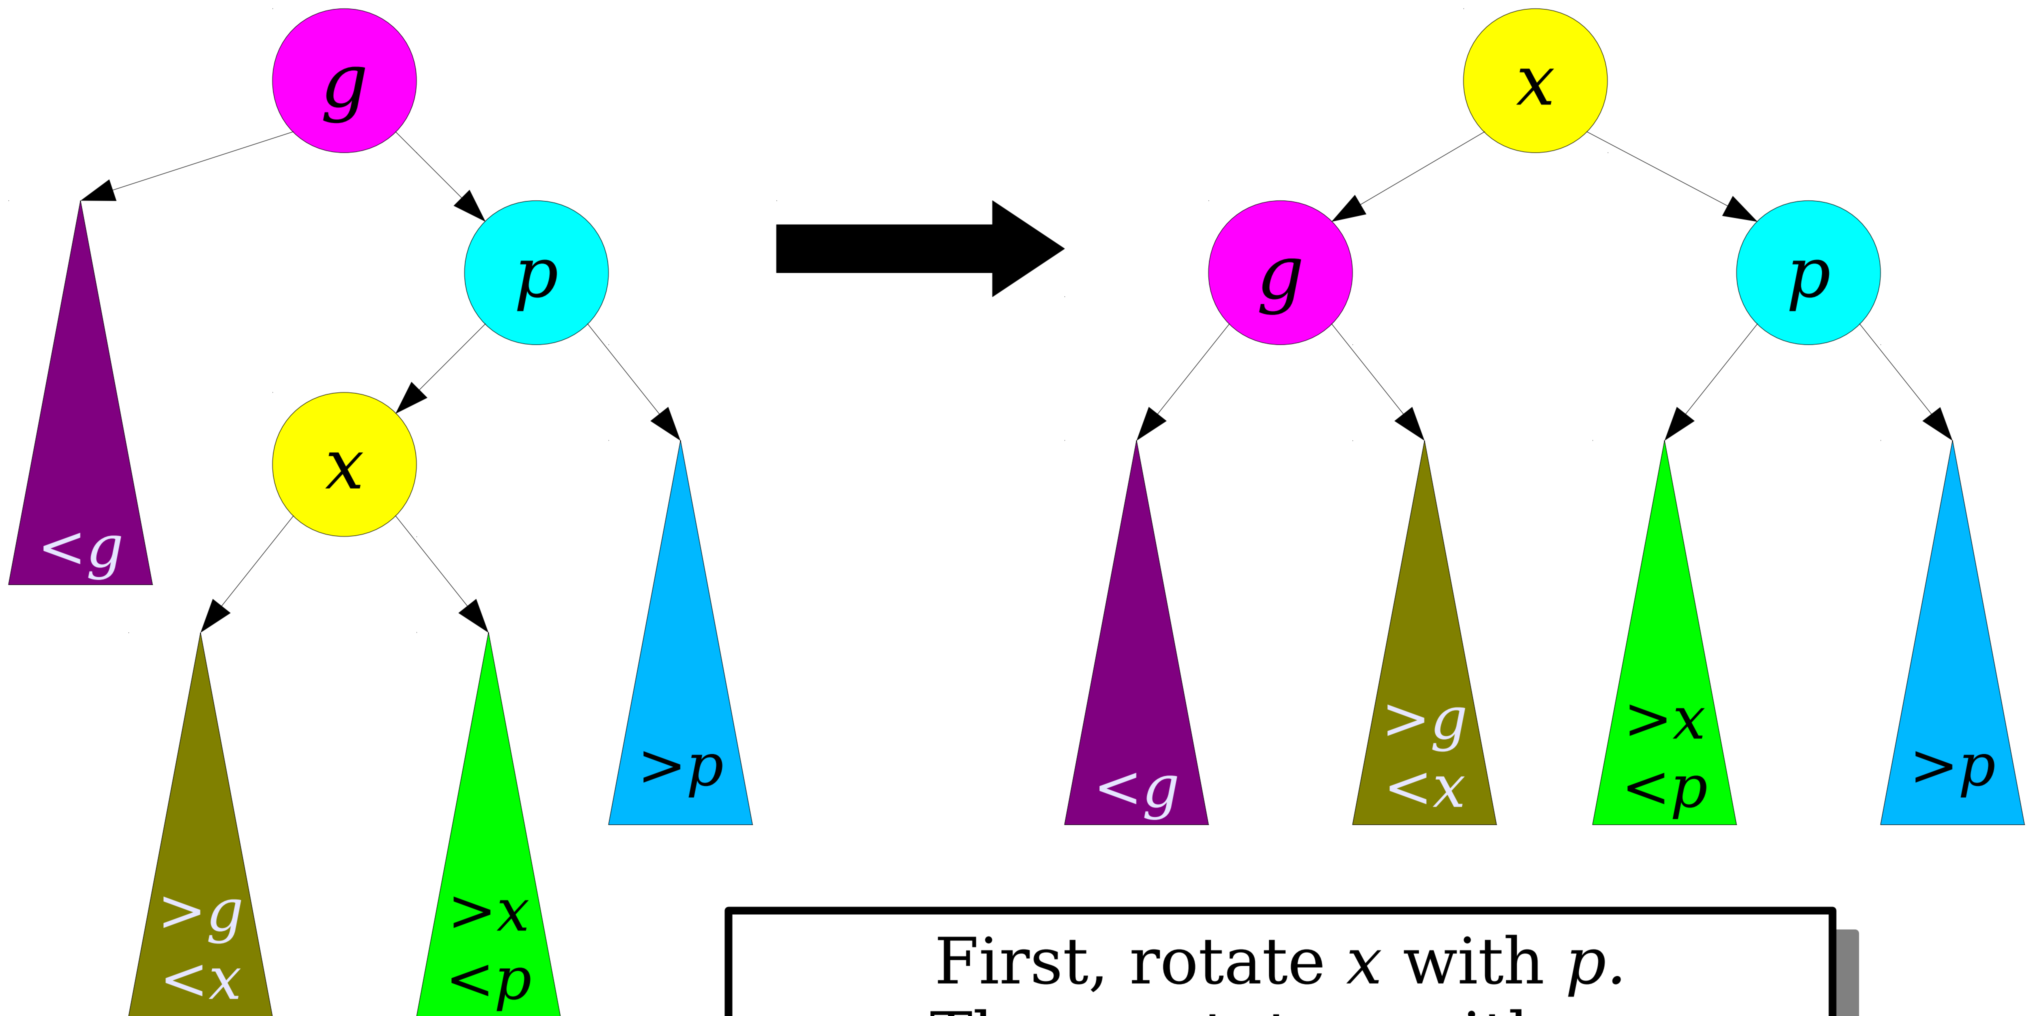

First, rotate x with p.
Then, rotate x with g.
Continue moving x up the tree.

Case 2: Zig-Zig

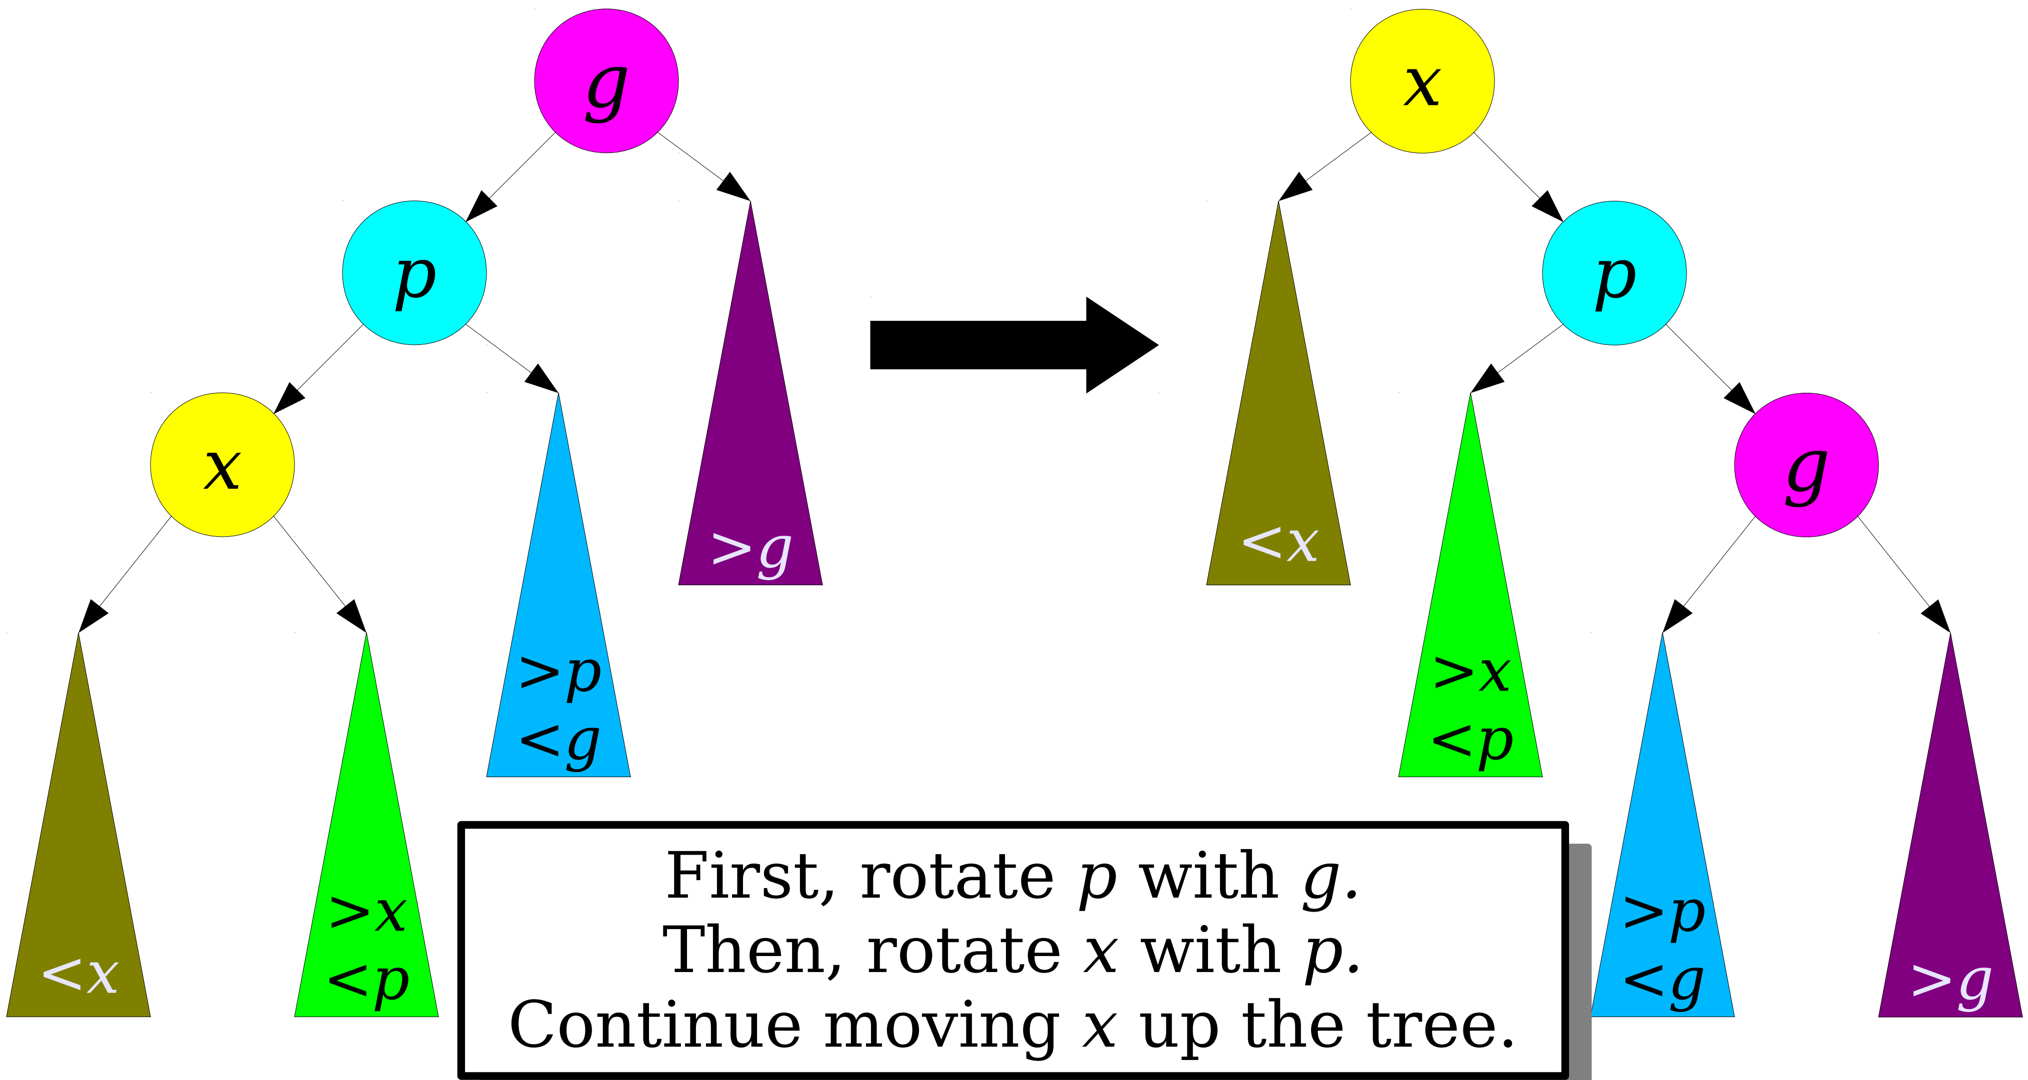

Case 3: Zig

(Assume r is the tree root)

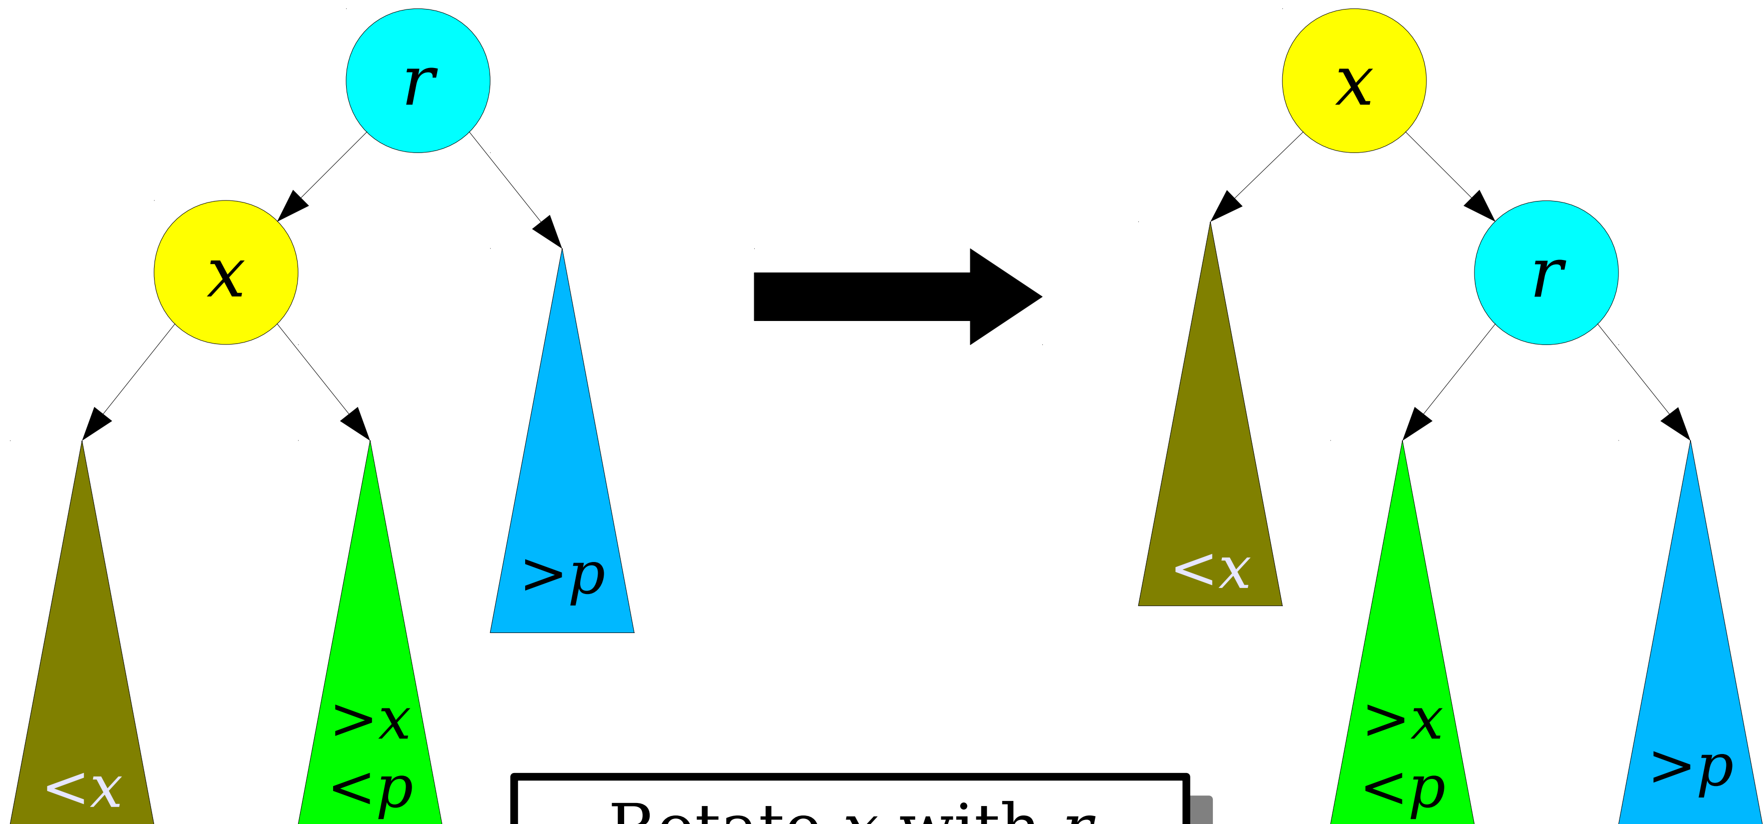

Rotate x with r
x is now the root.

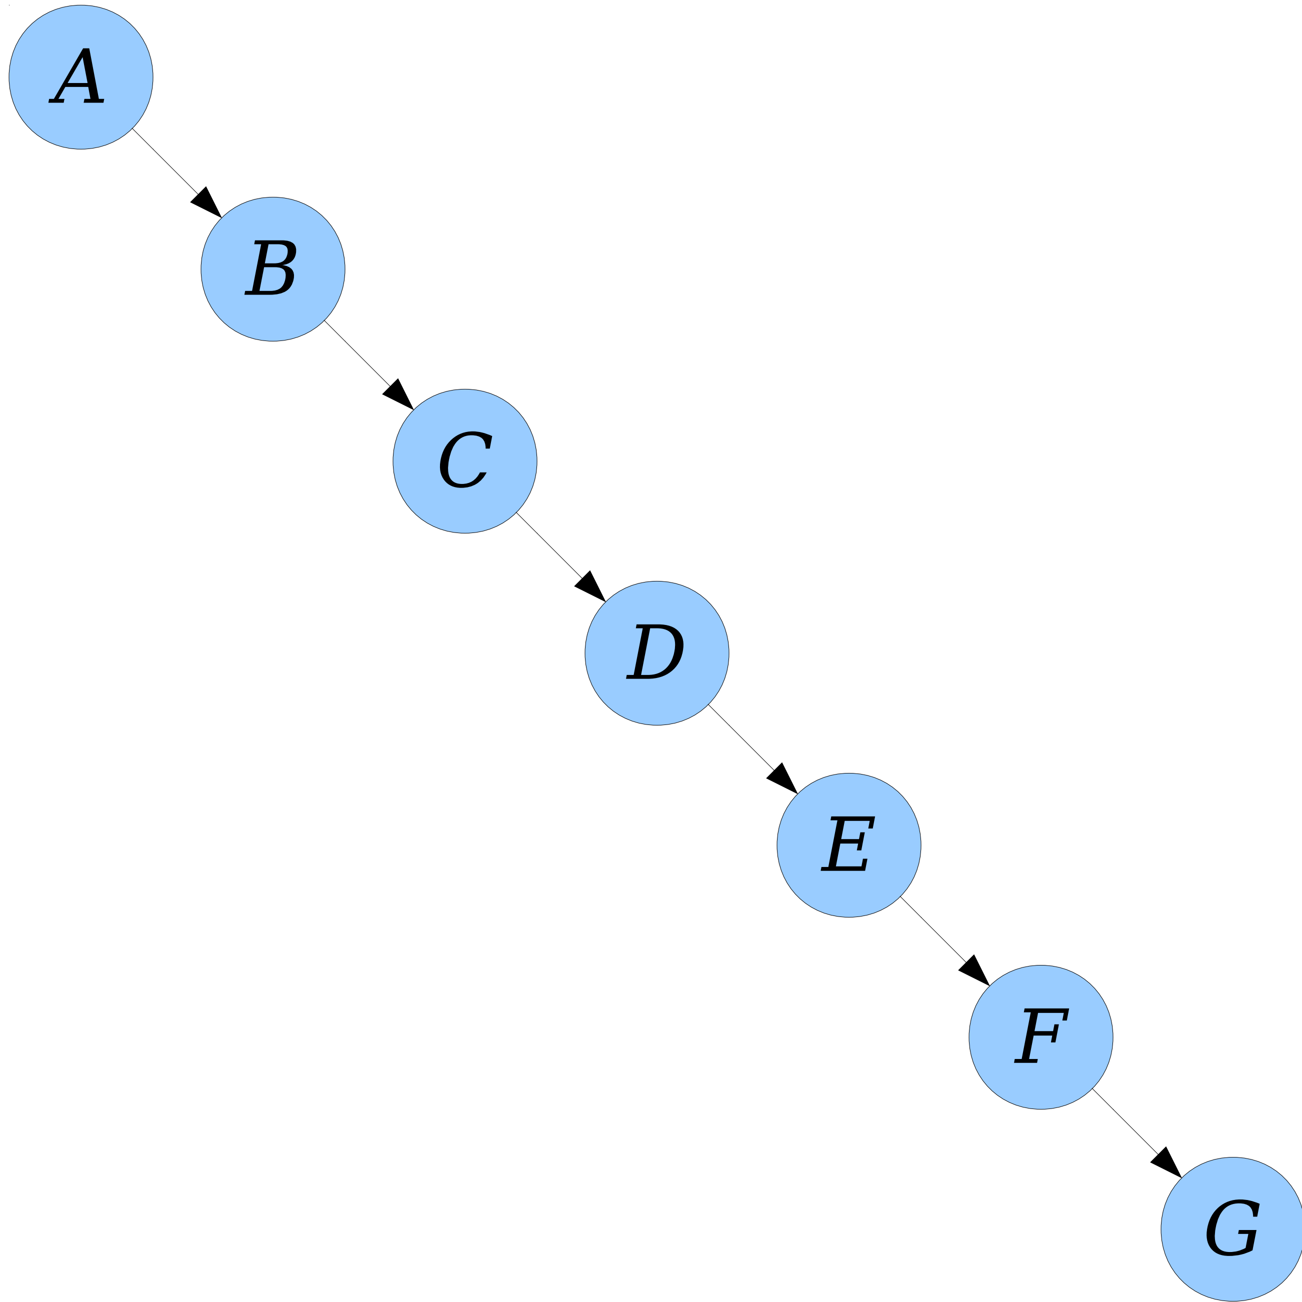
A
B
C
D
E
F
G

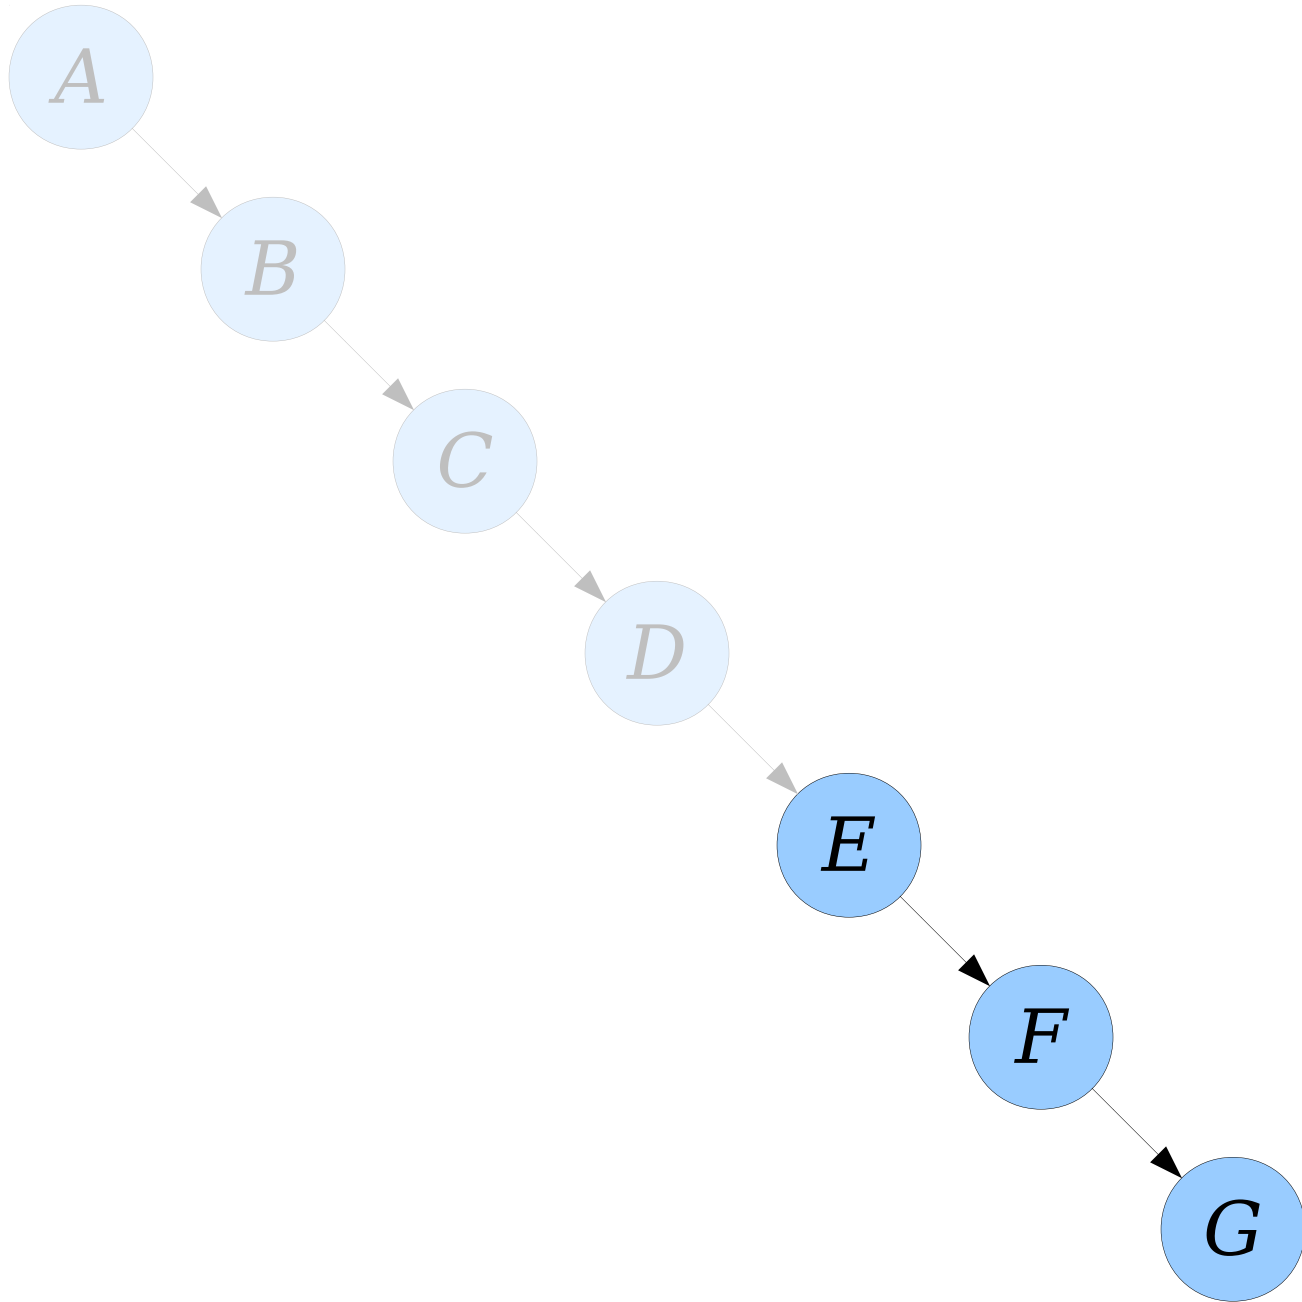

A
B
C
D
E
F
G

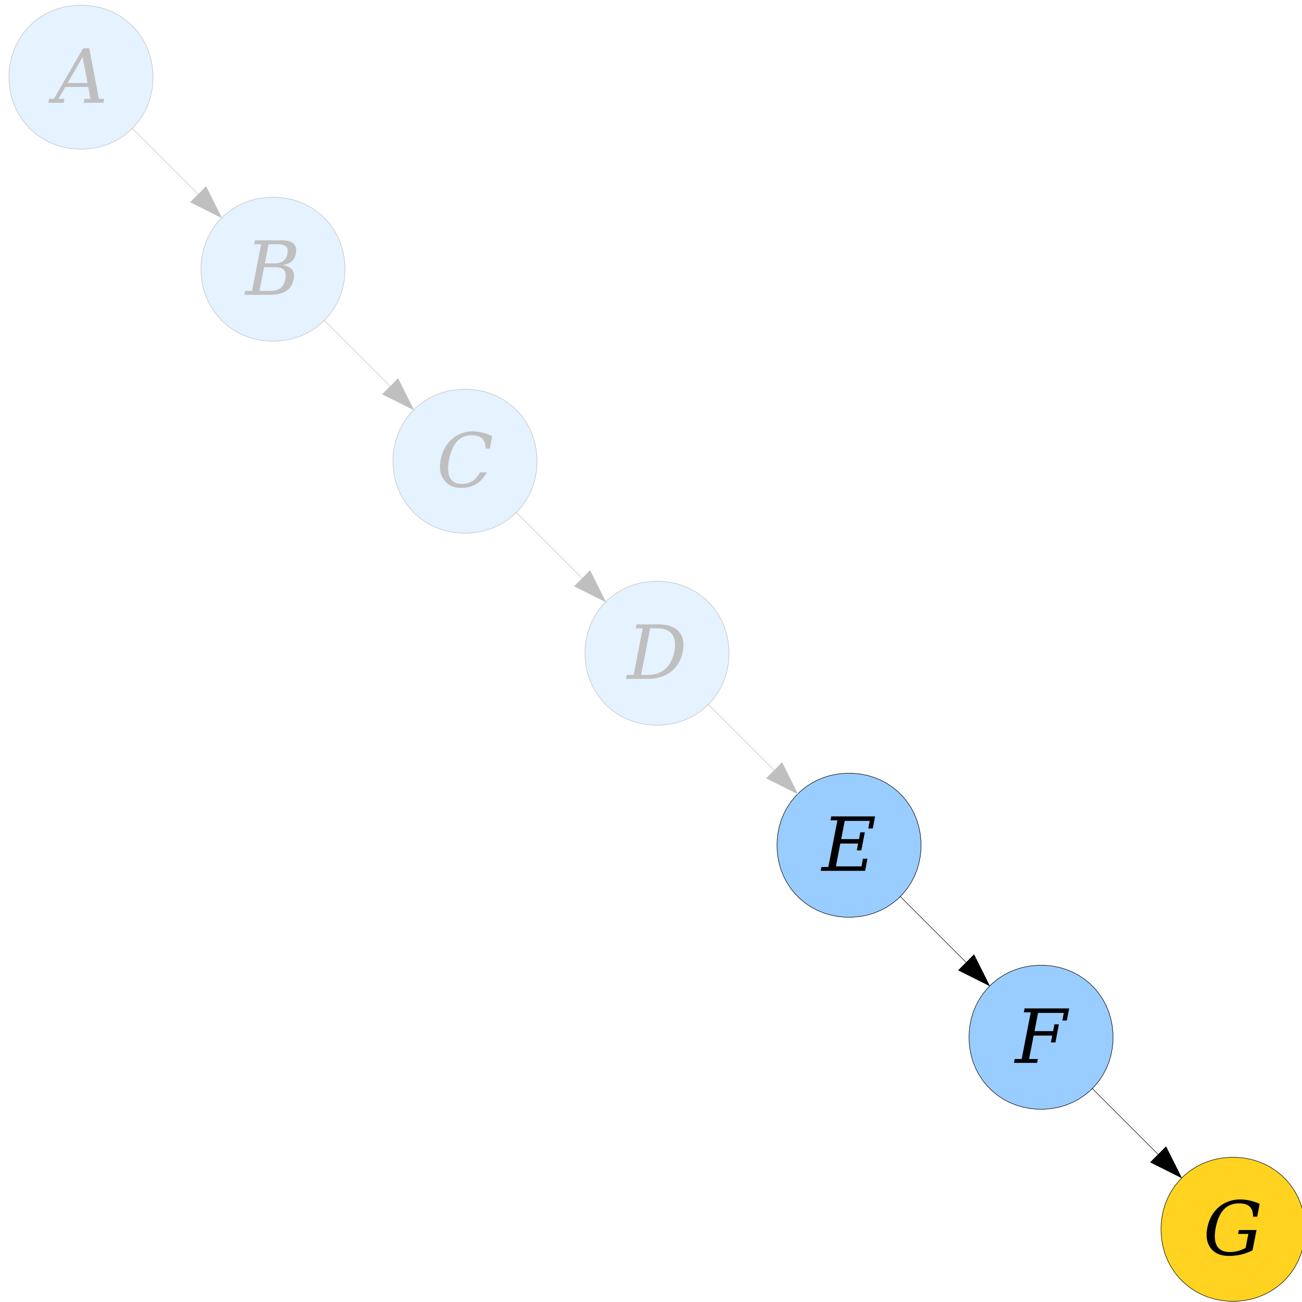
A
B
C
D
E
F
G

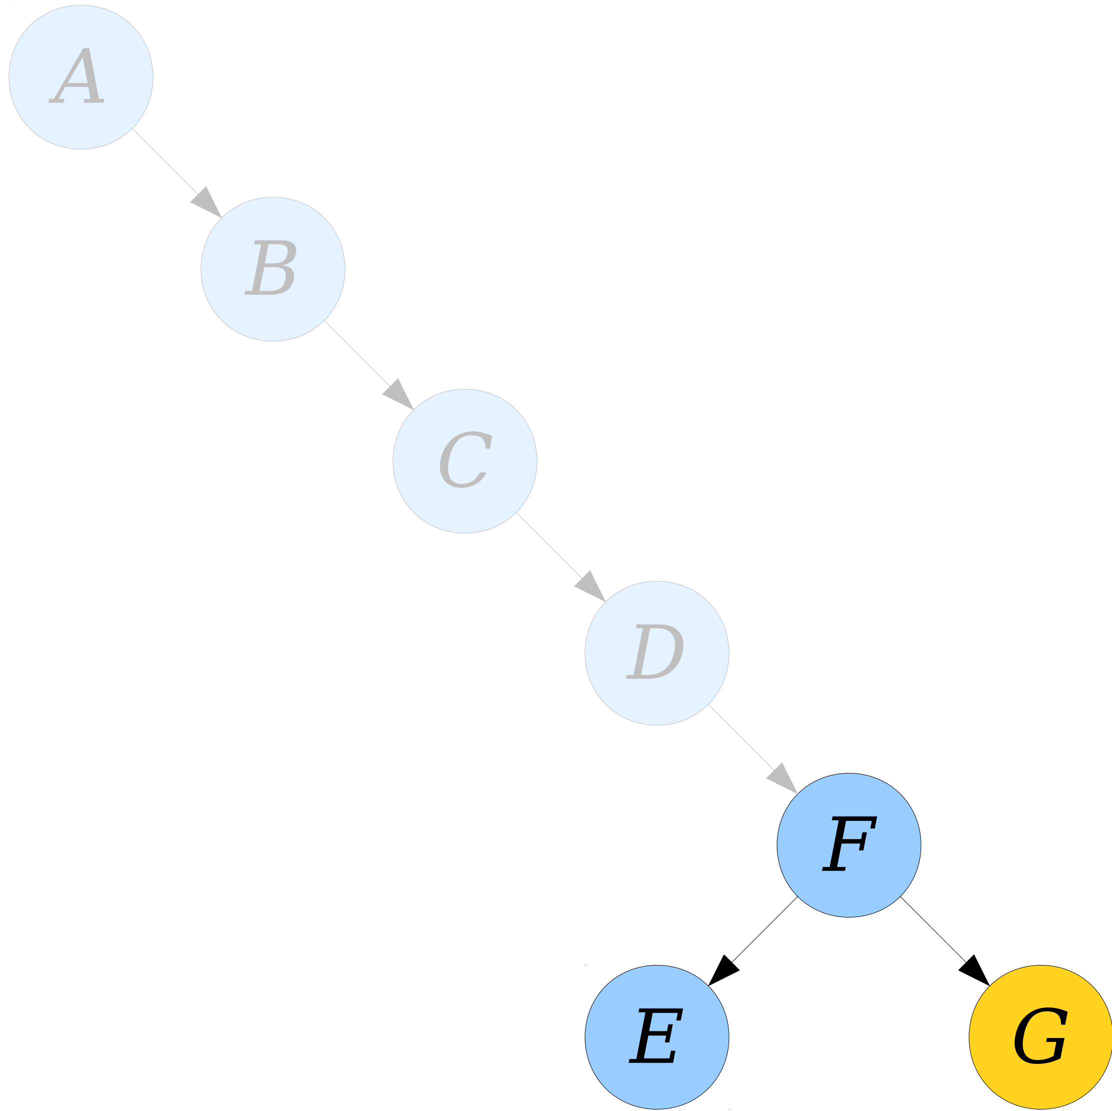
A
B
C
D
F
E
G

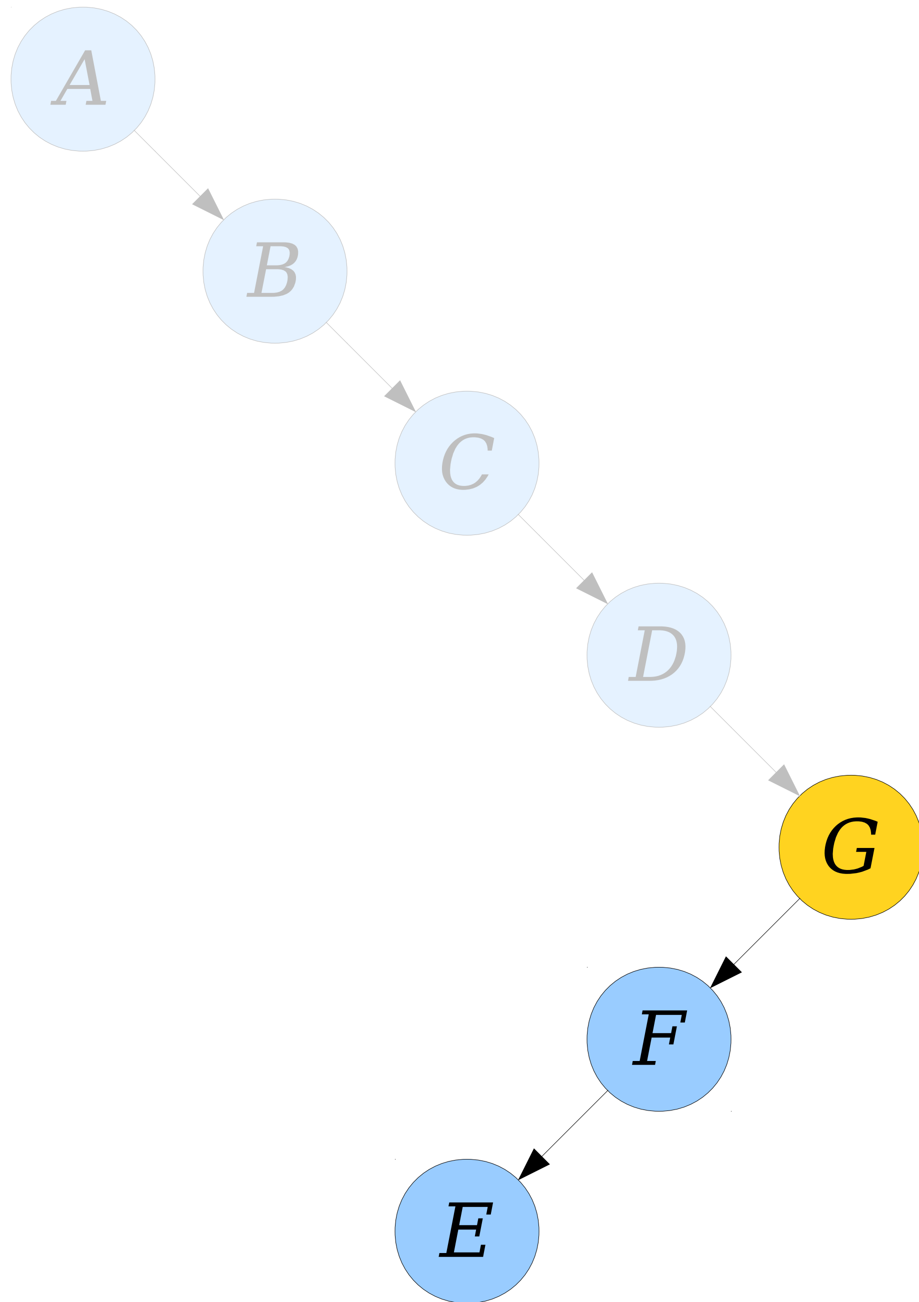

A
B
C
D
G
F
E

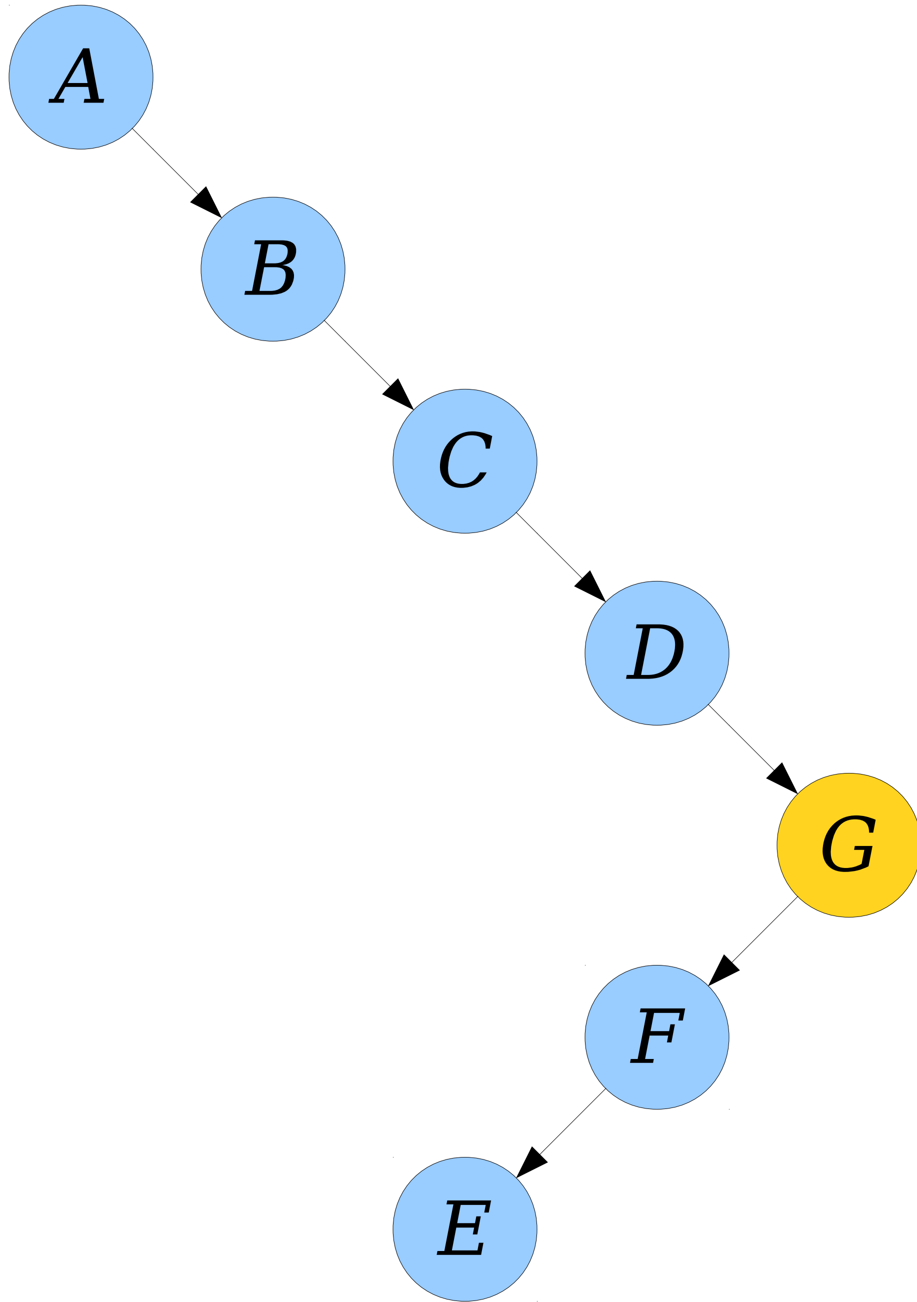
A
B
C
D
G
F
E

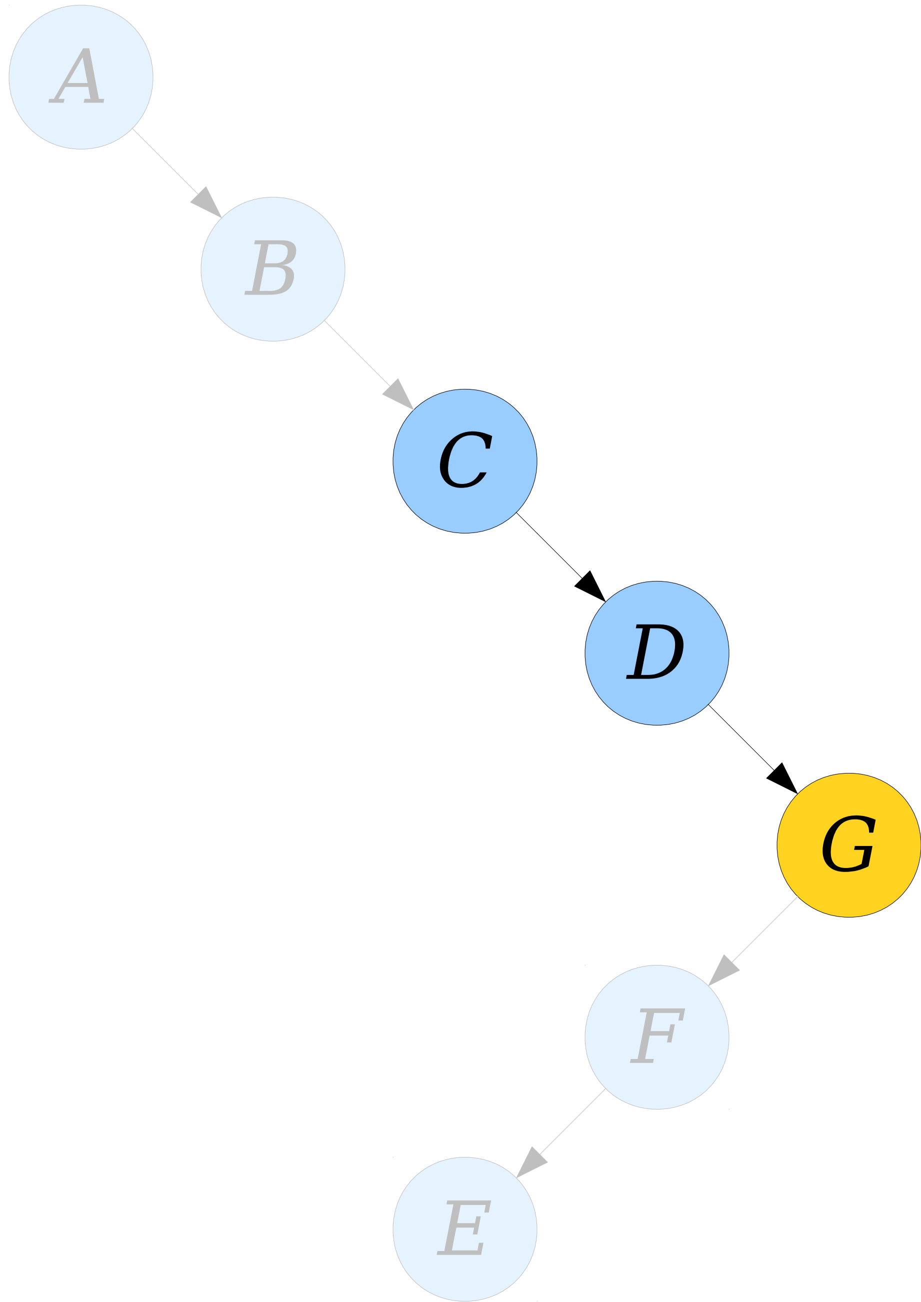
A
B
C
D
G
F
E

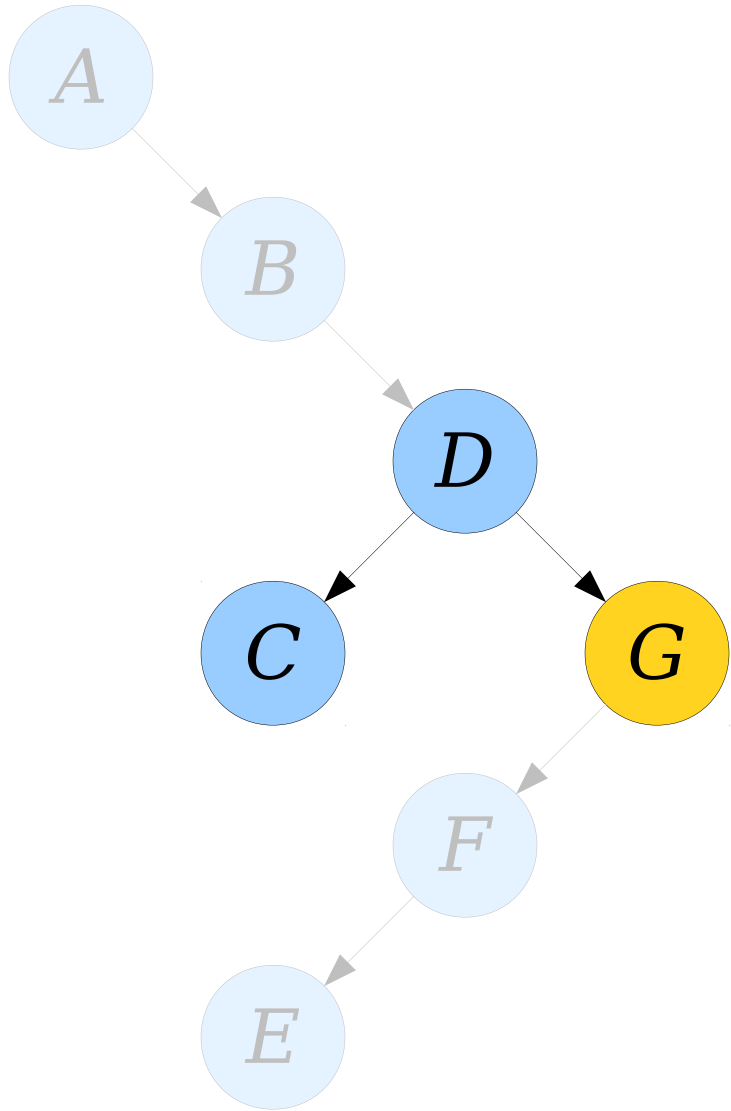

A
B
D
C
G
F
E

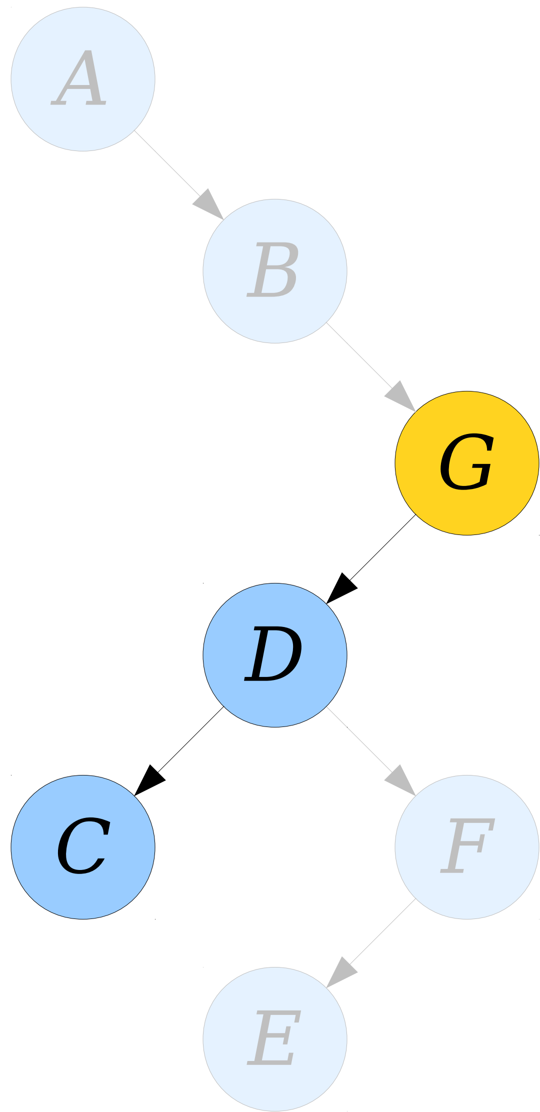
A
B
G
D
C
F
E

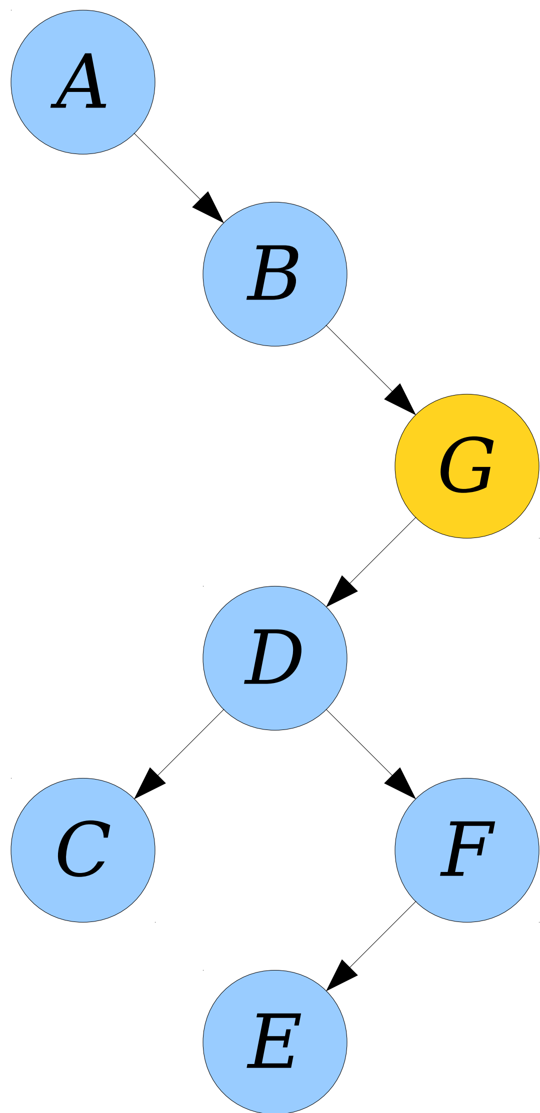

A
B
G
D
C
F
E

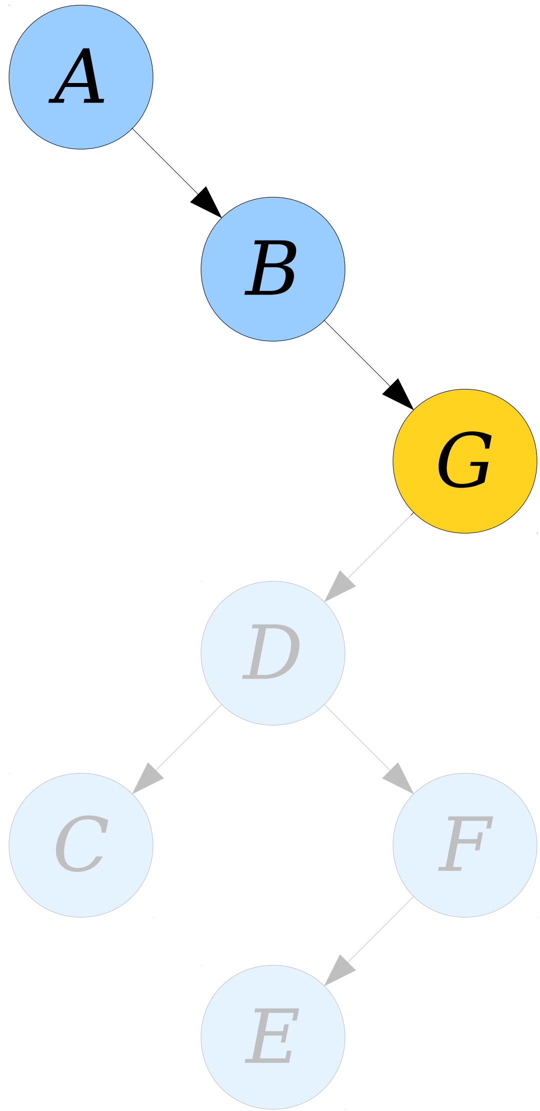
A
B
G
D
C
F
E

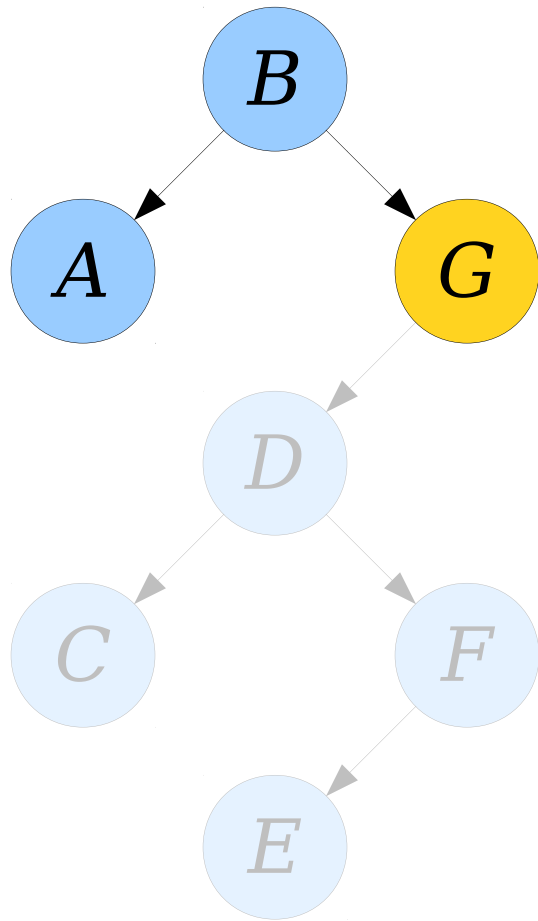
B
A
G
D
C
F
E

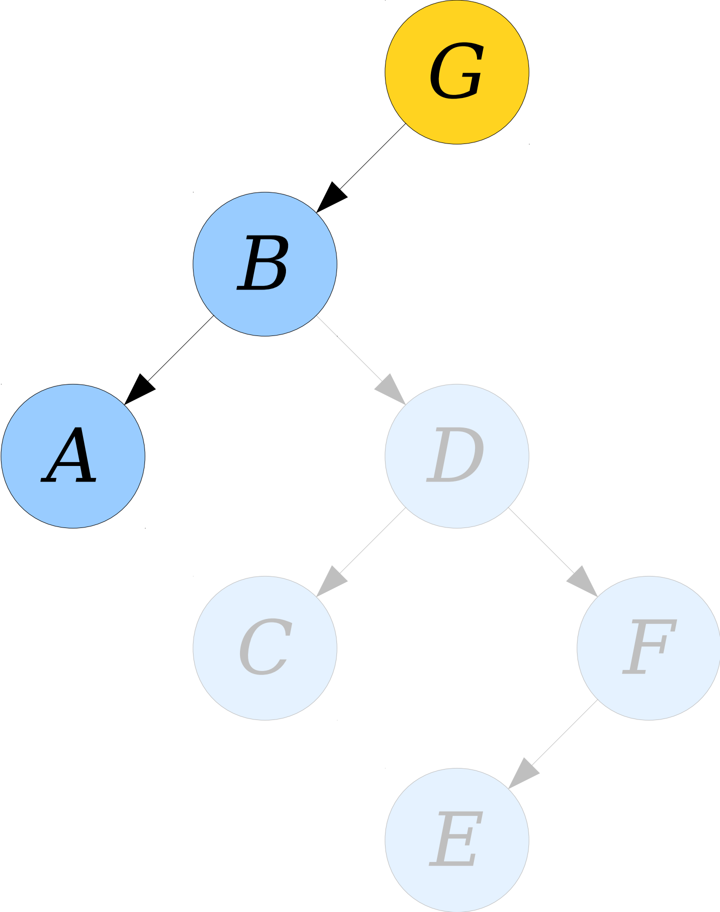
G
B
A
D
C
F
E

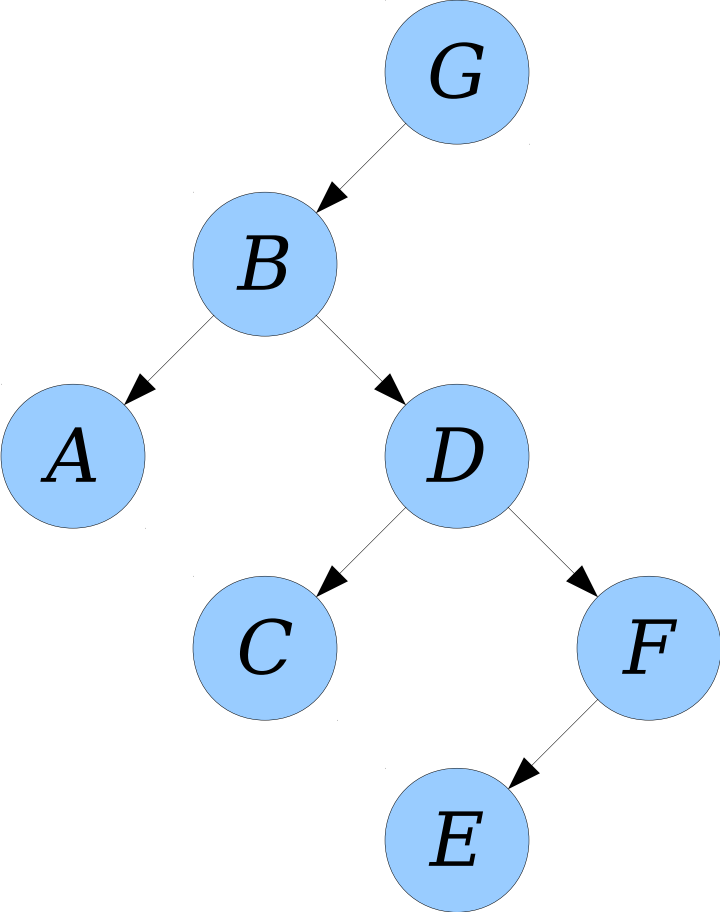
G
B
A
D
C
F
E

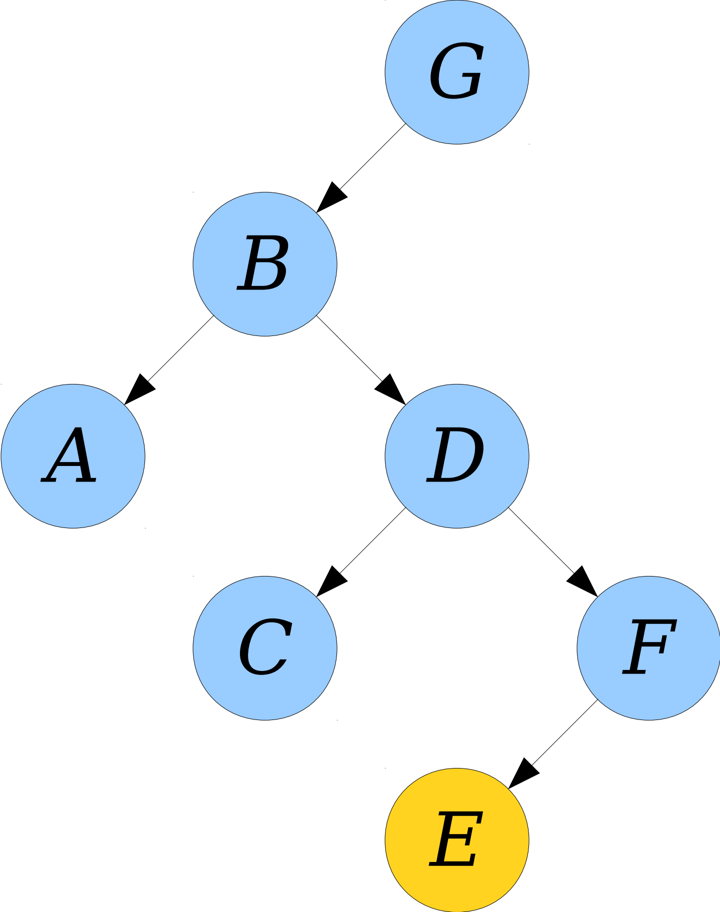
G
B
A
D
C
F
E

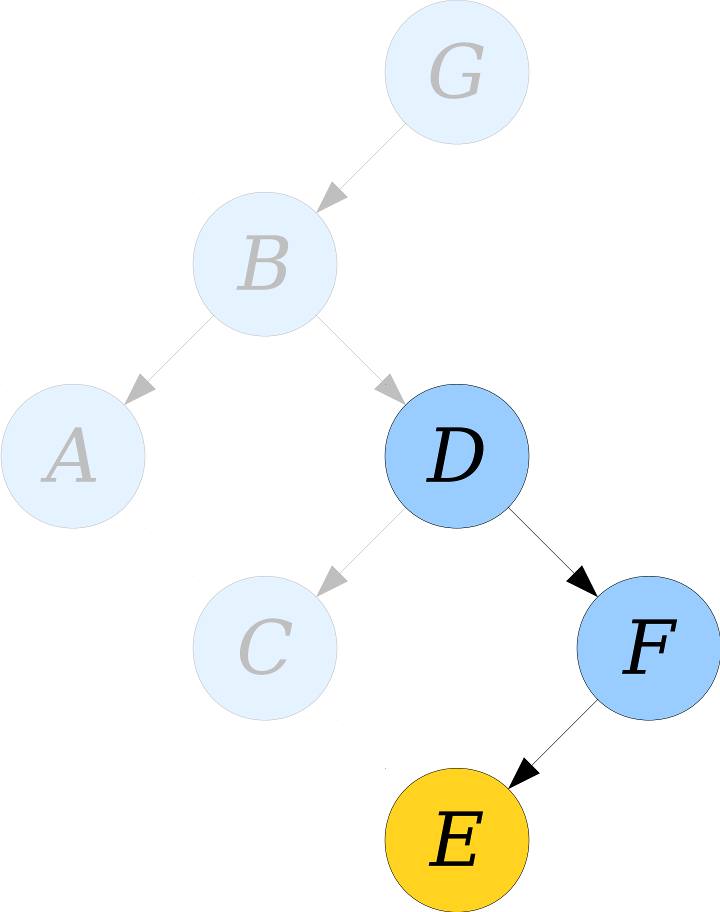
G
B
A
D
C
F
E

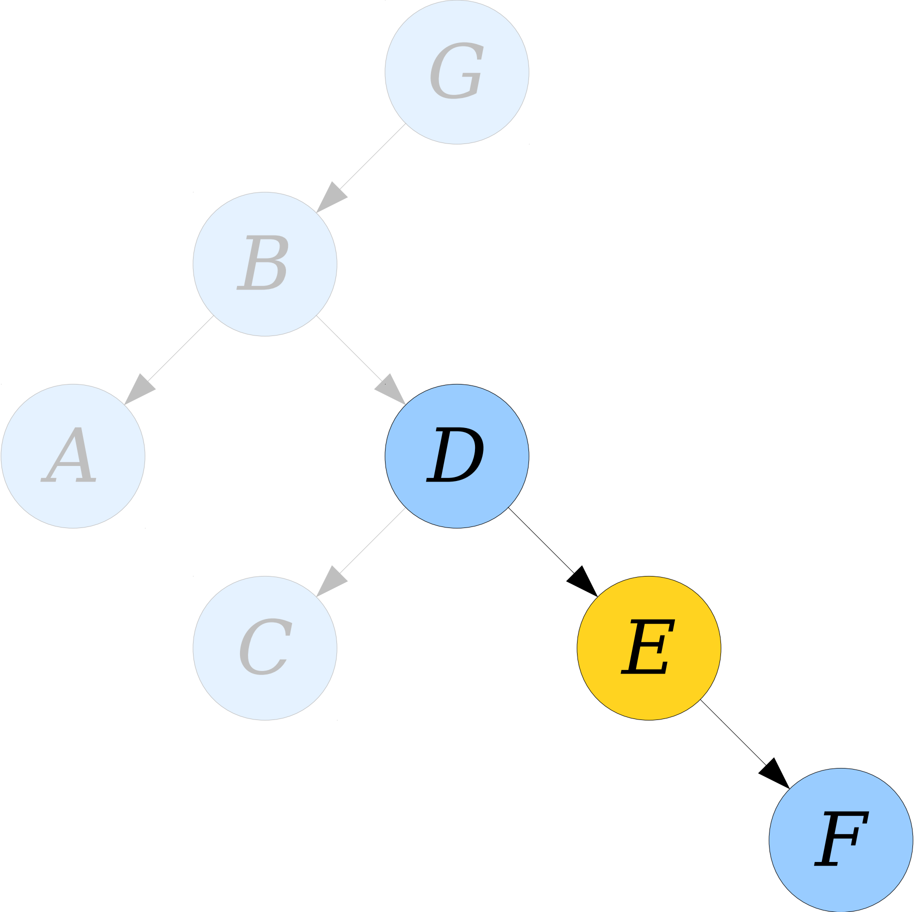
G
B
A
D
C
E
F

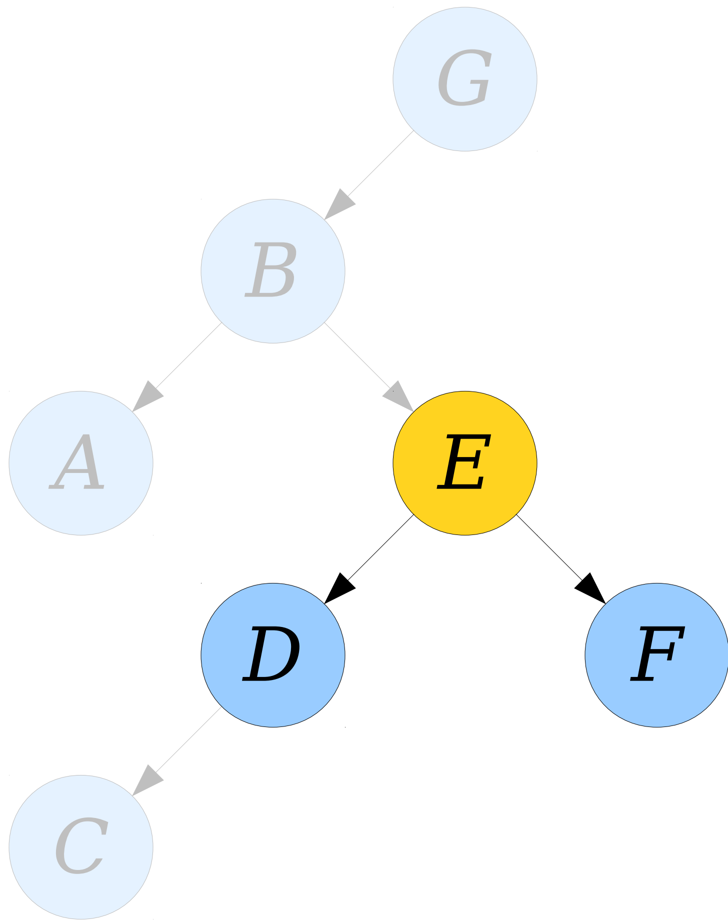
G
B
A
E
D
F
C

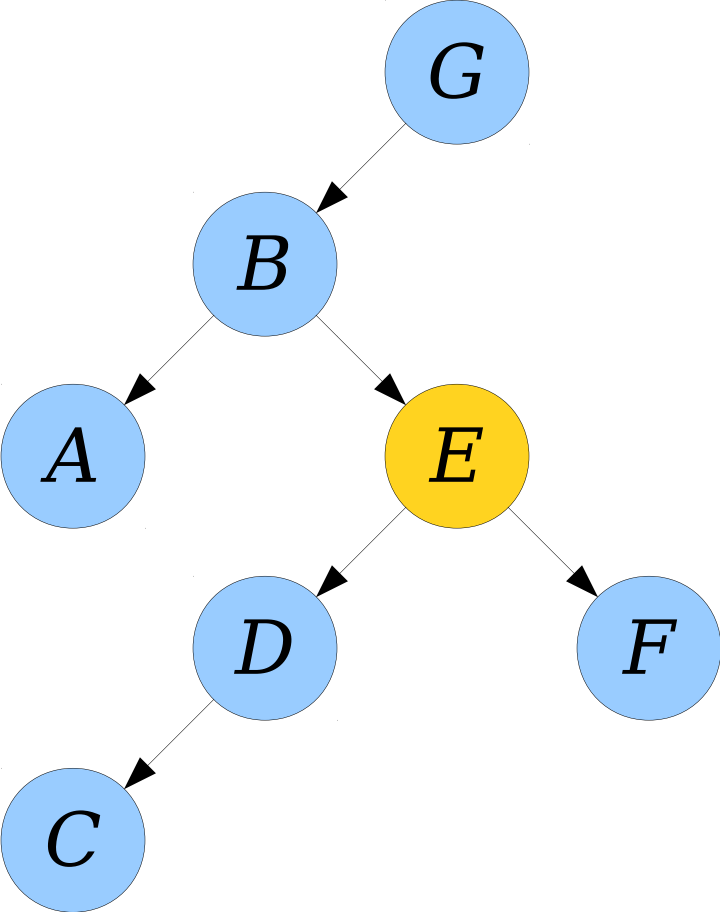
G
B
A
E
D
F
C

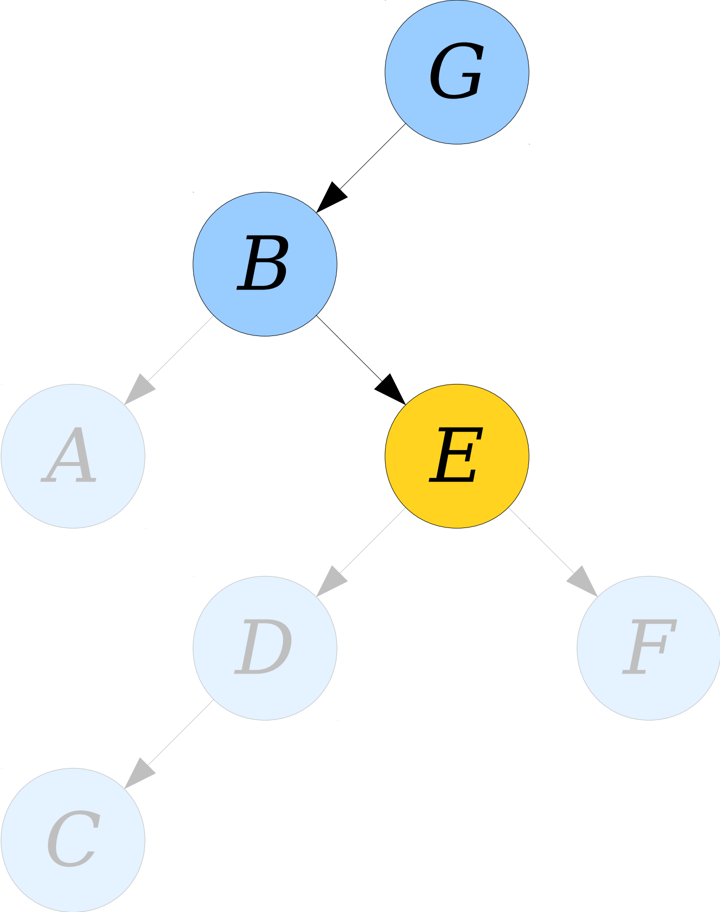
G
B
A
E
D
F
C

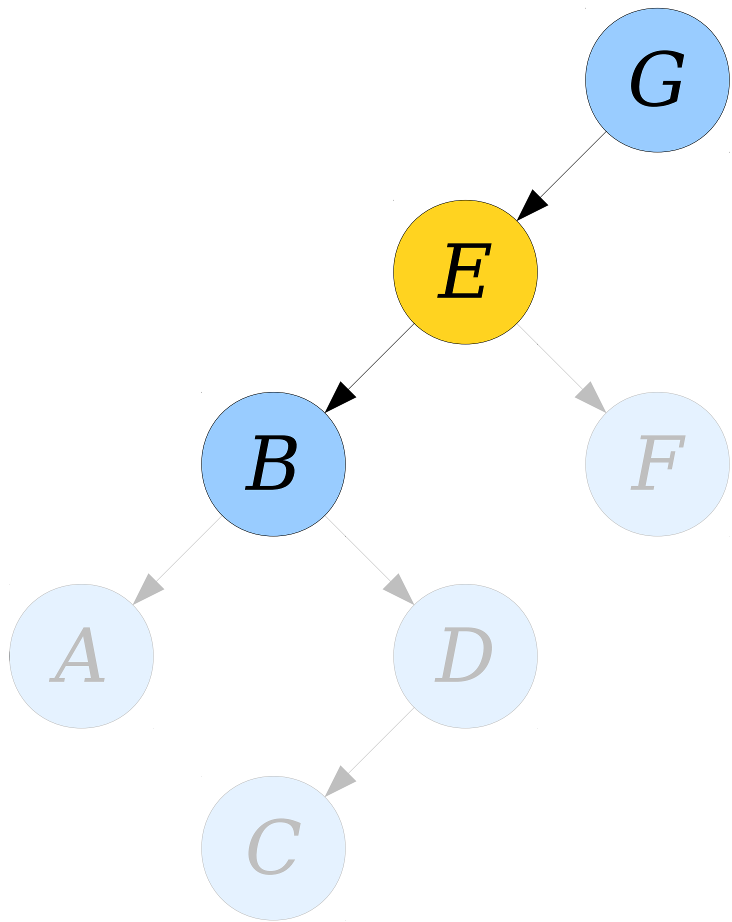
G
E
B
F
A
D
C

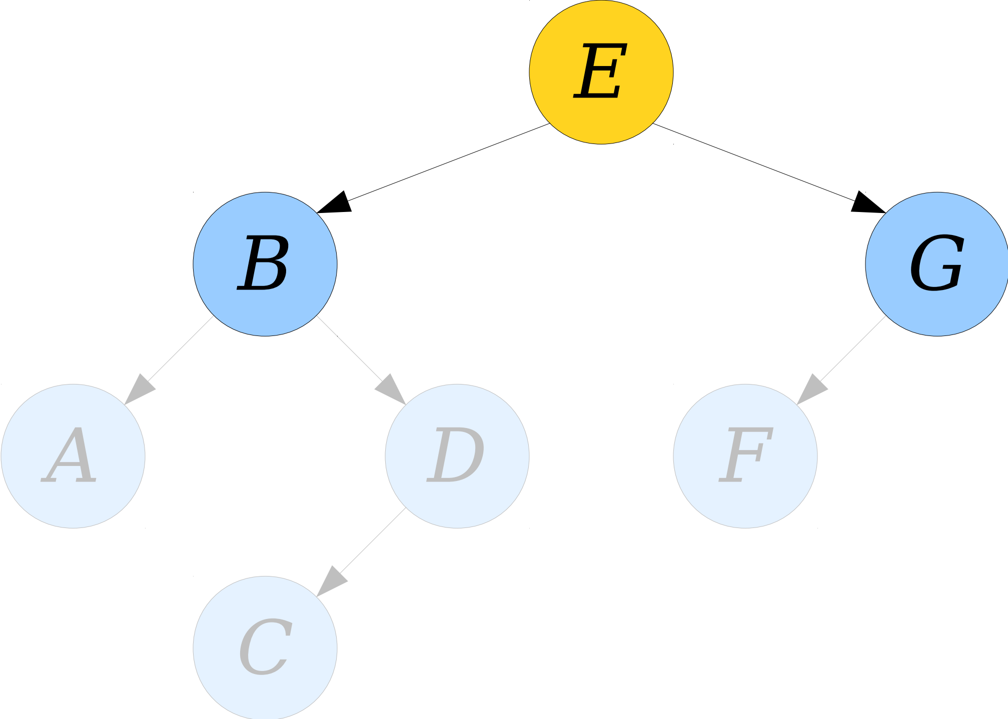
E
B
G
A
D
F
C

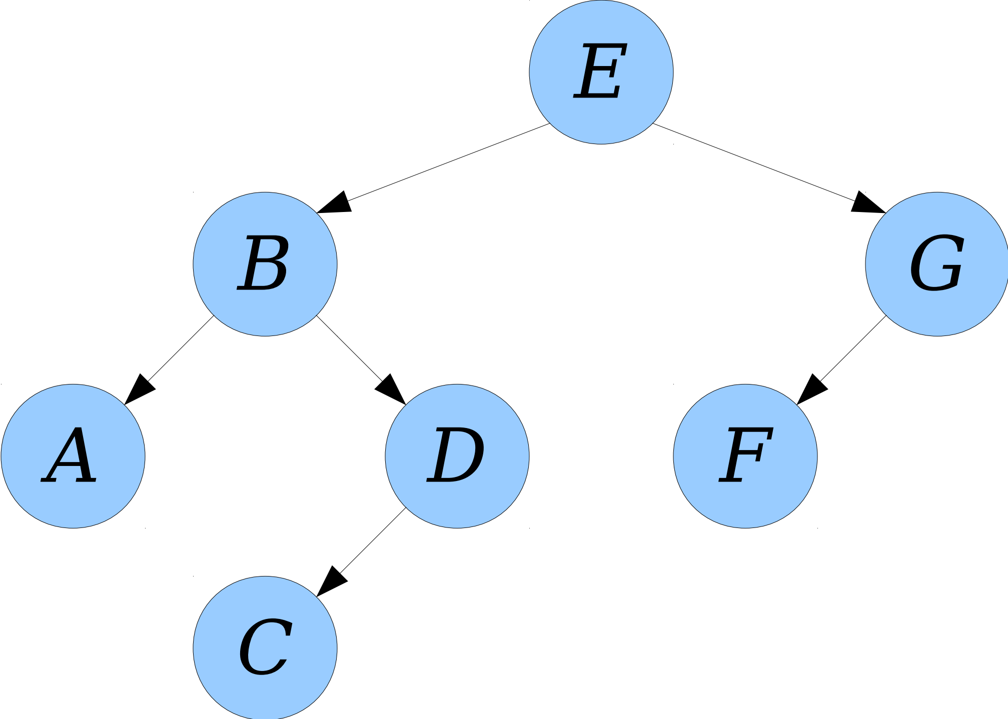
E
B
G
A
D
F
C

Splaying, Empirically

- After a few splays, we went from a totally degenerate tree to a reasonably-balanced tree.
- Splaying nodes that are deep in the tree tends to correct the tree shape.
- Why is this?
- Is this a coincidence?

Why Splaying Works

- ***Claim:*** After doing a splay at x, the average depths of any nodes on the access path to x are halved.
- This helps eliminate long access paths, making lookups of elements near x faster in the future.

The average depth of *x*, *p*, and *g* is unchanged.

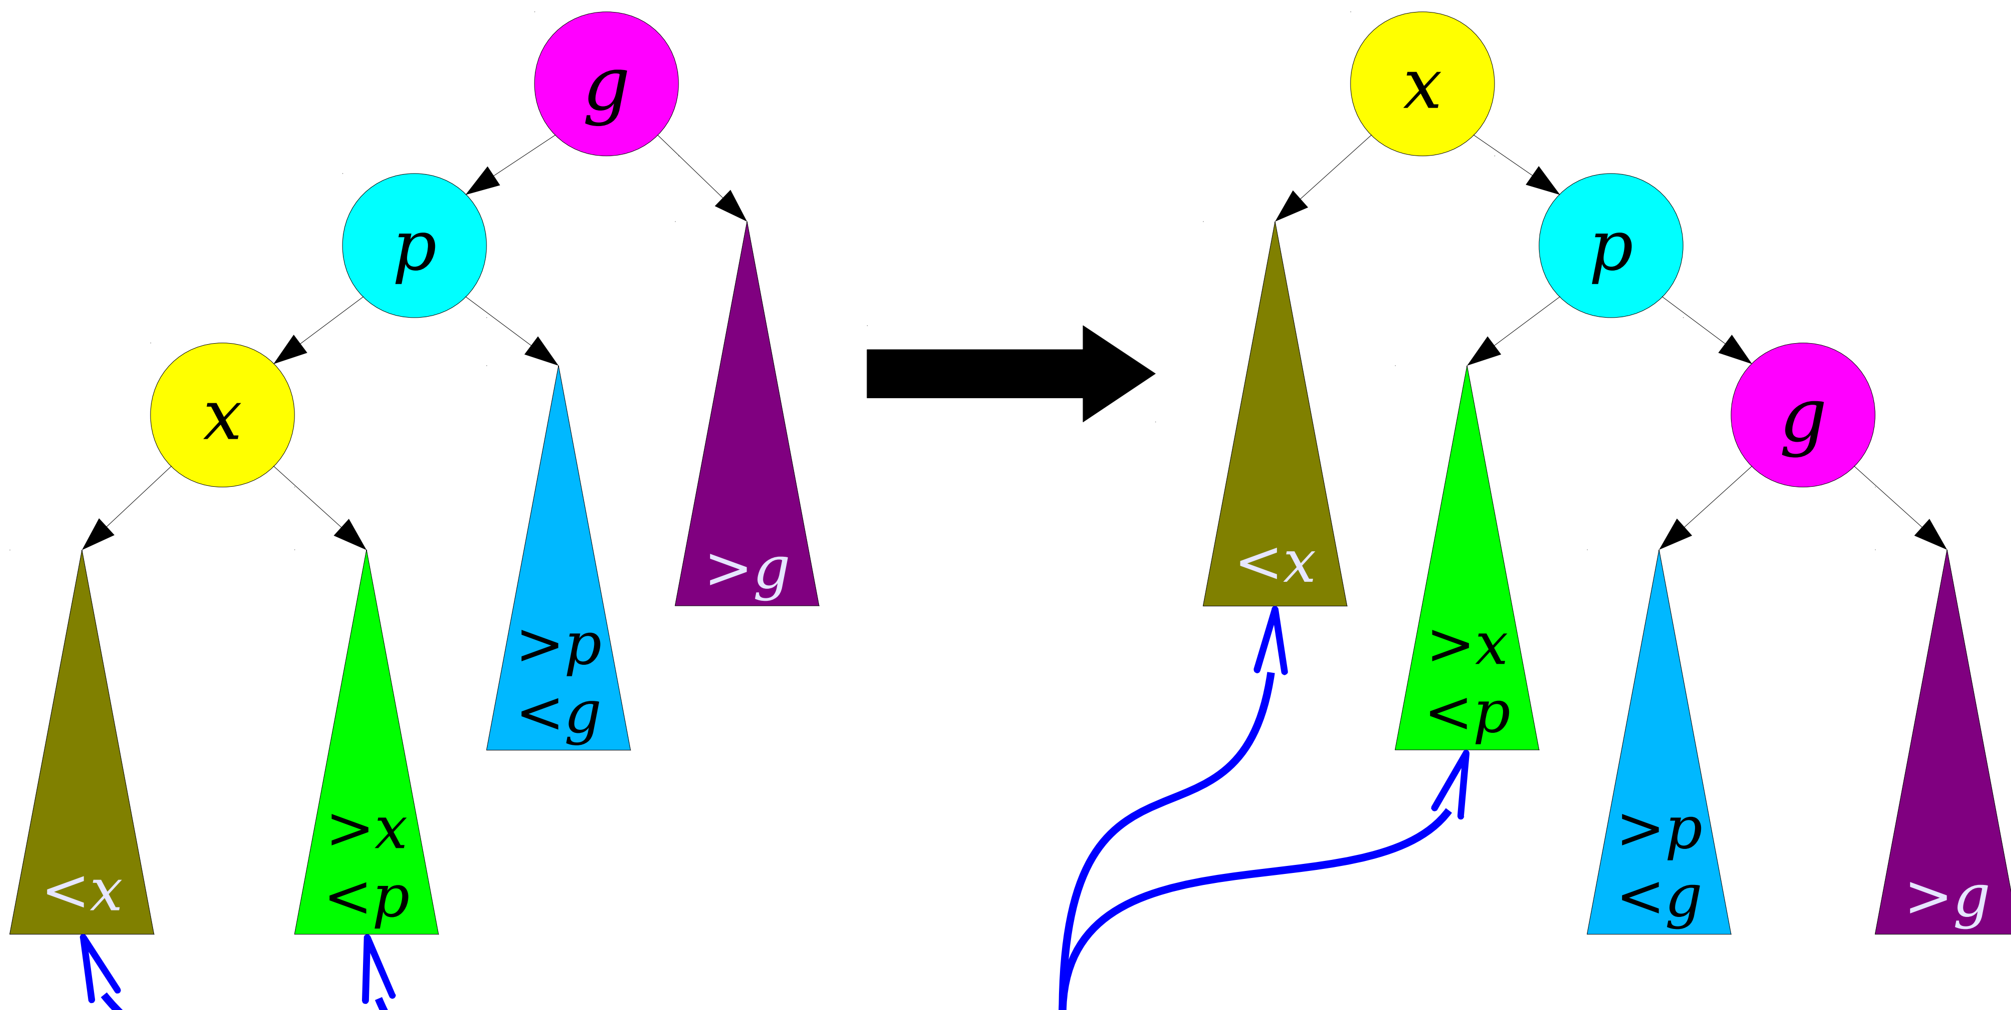

These subtrees decrease in height by one or two.

The average height of *x*, *p*, and *g* decreases by $^{1}/_{3}$.

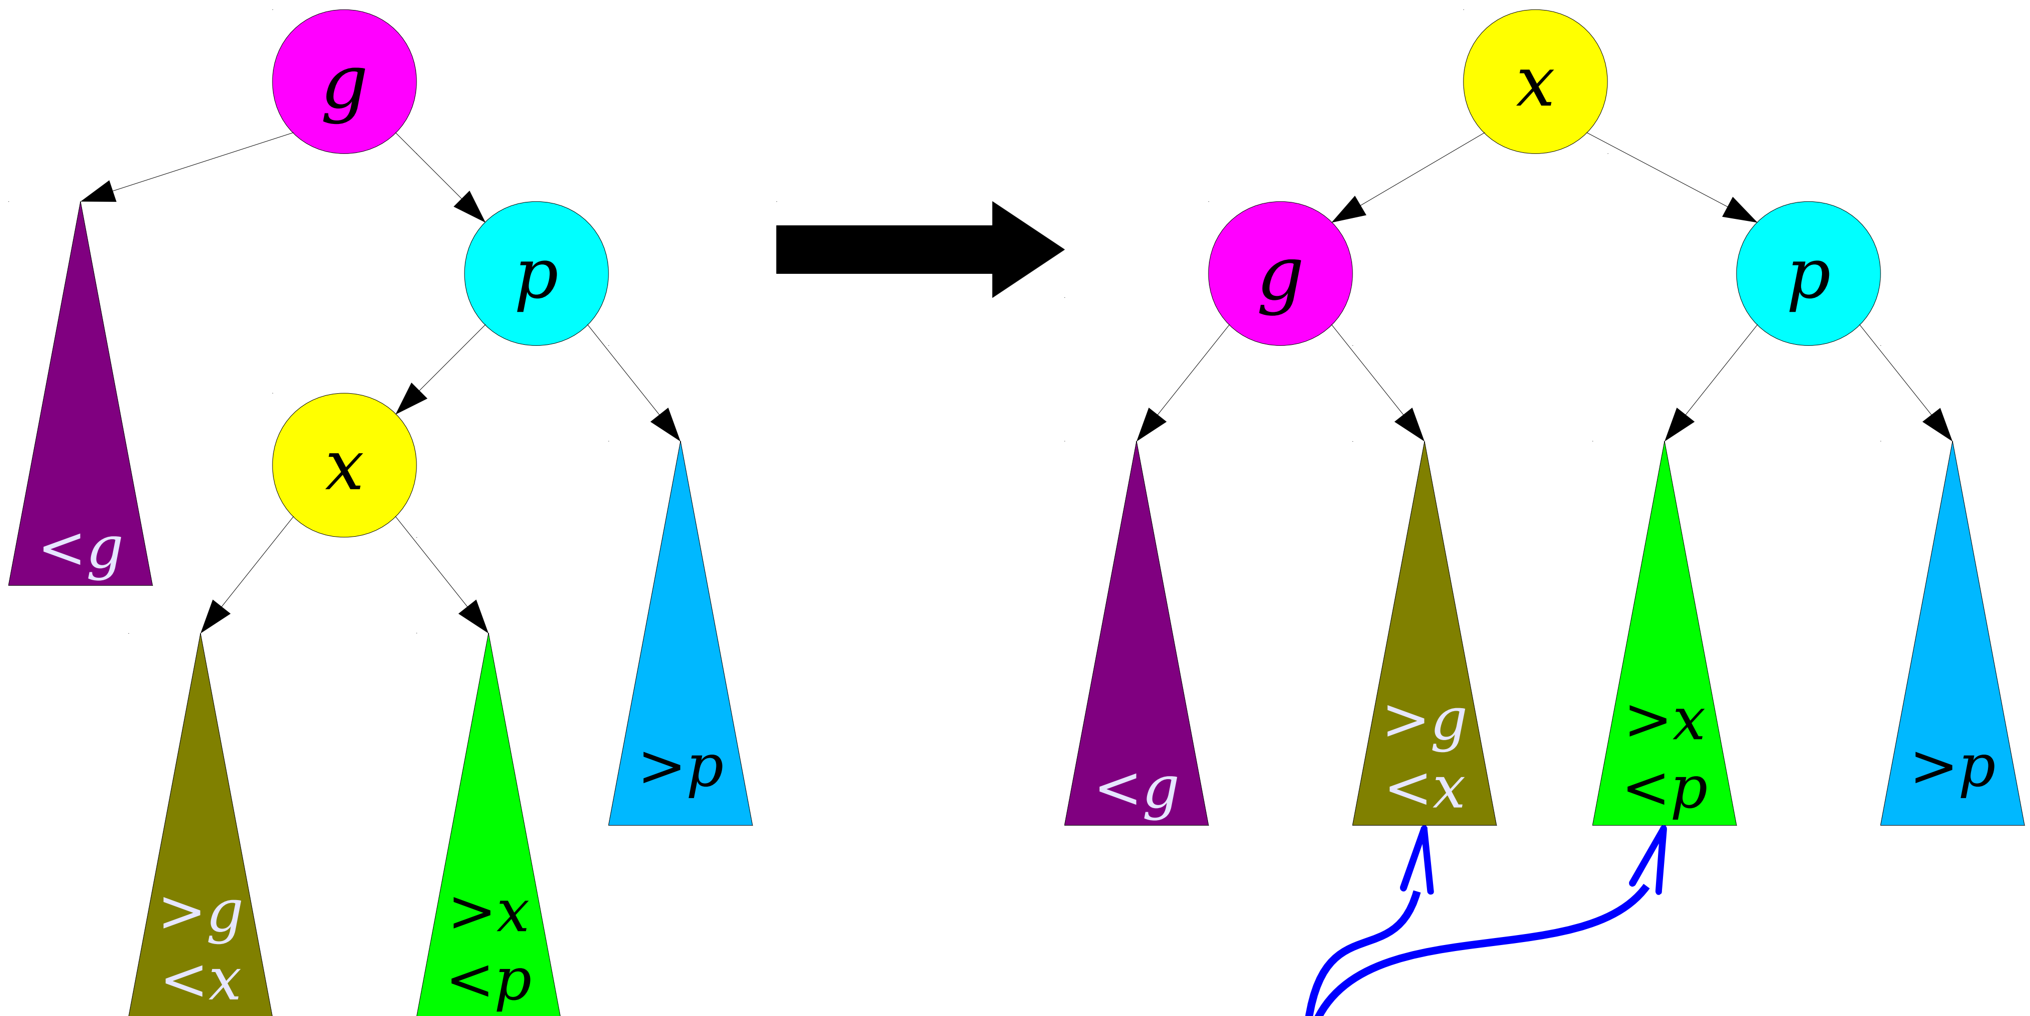

These subtrees have their height decreased by one.

There is no net change in the height of x or r.

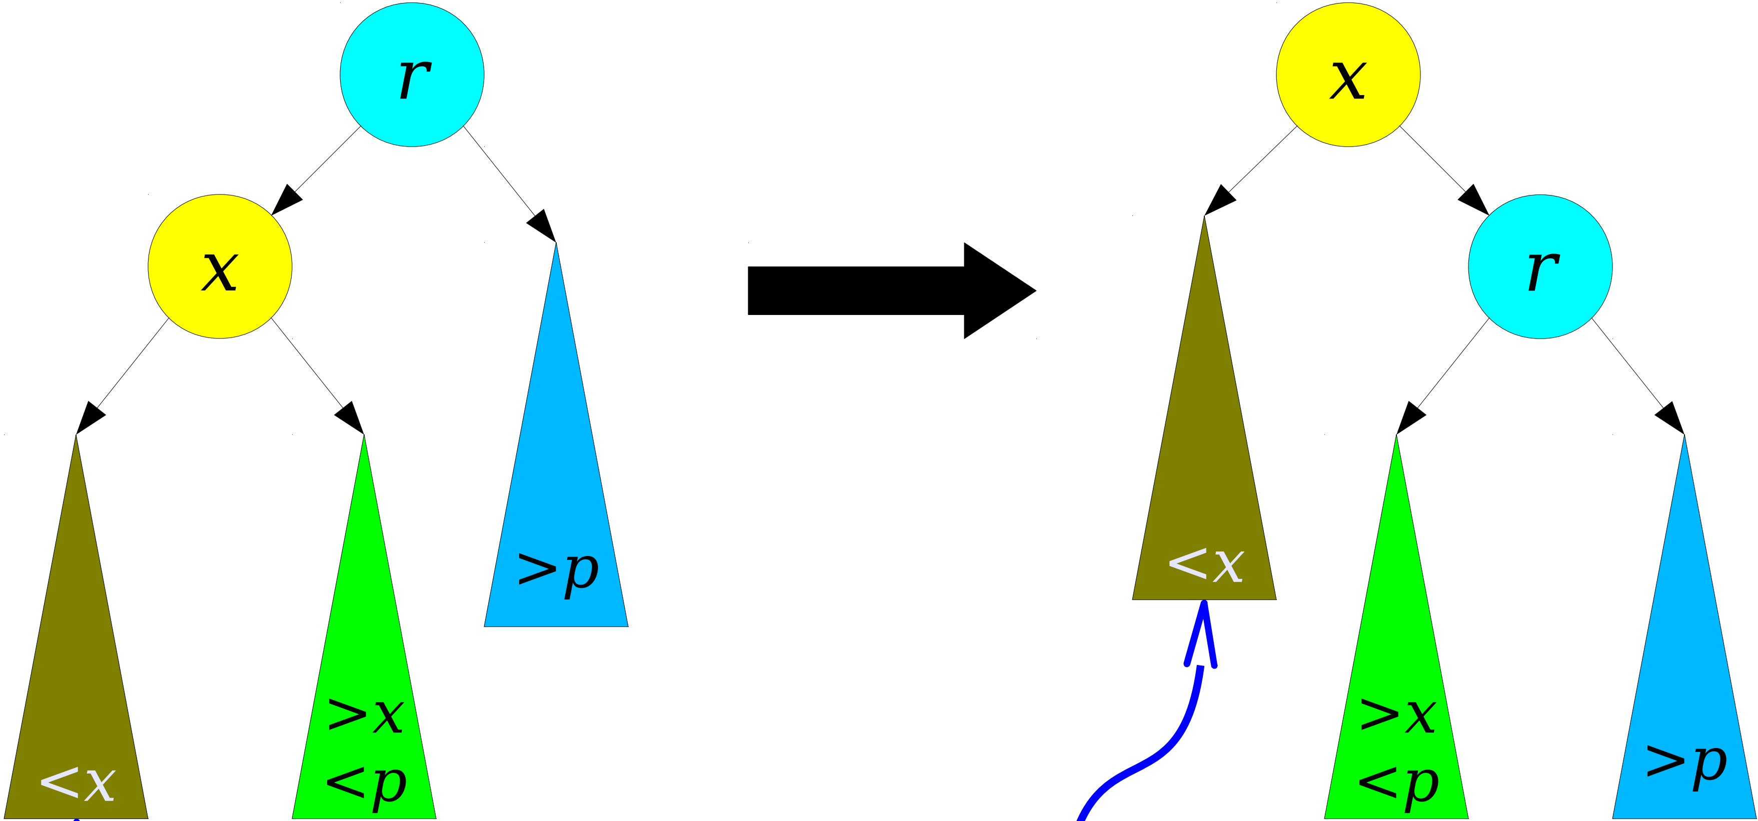

The nodes in this subtree have their height decreased by one.

An Intuition for Splaying

- Each rotation done only slightly penalizes each other part of the tree (say, adding +1 or +2 depth).
- Each splay rapidly cuts down the height of each node on the access path.
- Slow growth in height, combined with rapid drop in height, is a hallmark of amortized efficiency.
- ***Claim:*** The original "rotate-to-root" idea from before doesn't do this, which partially explains why it's not a good strategy.

Rotate-to-Root

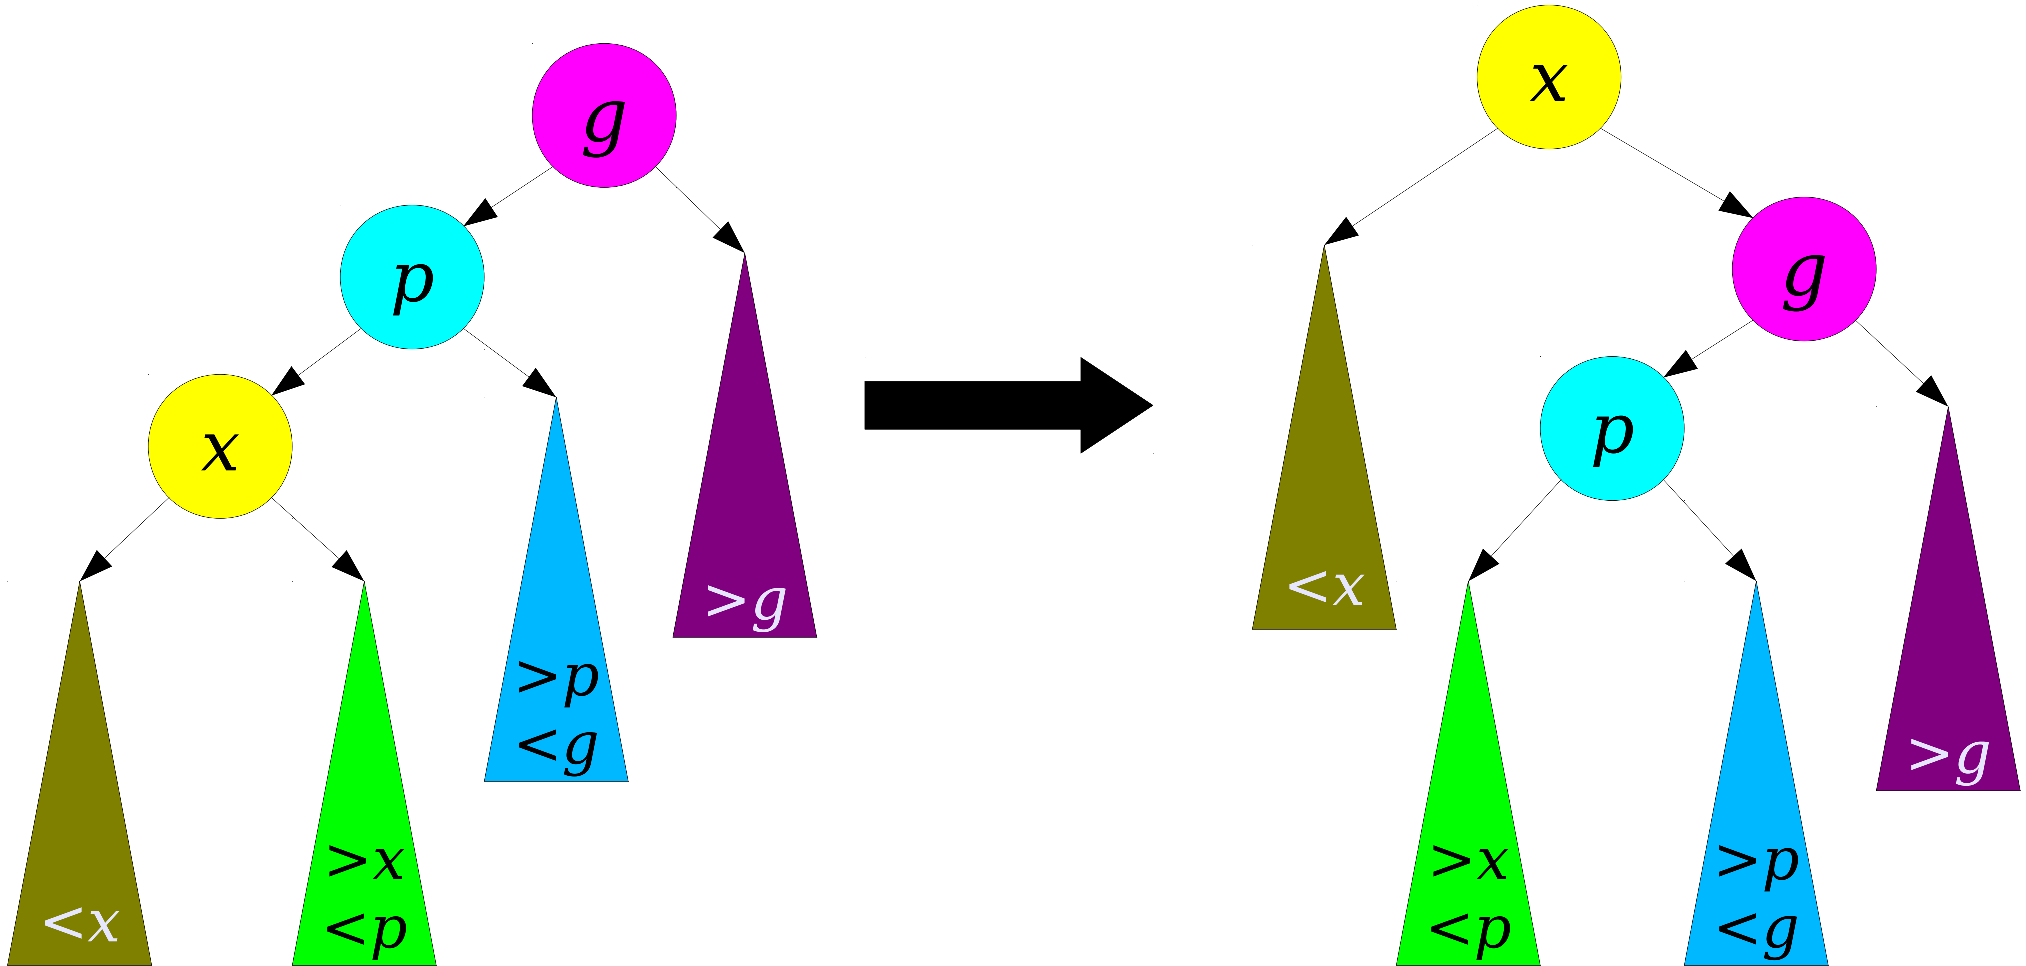

Why Rotate-to-Root Fails

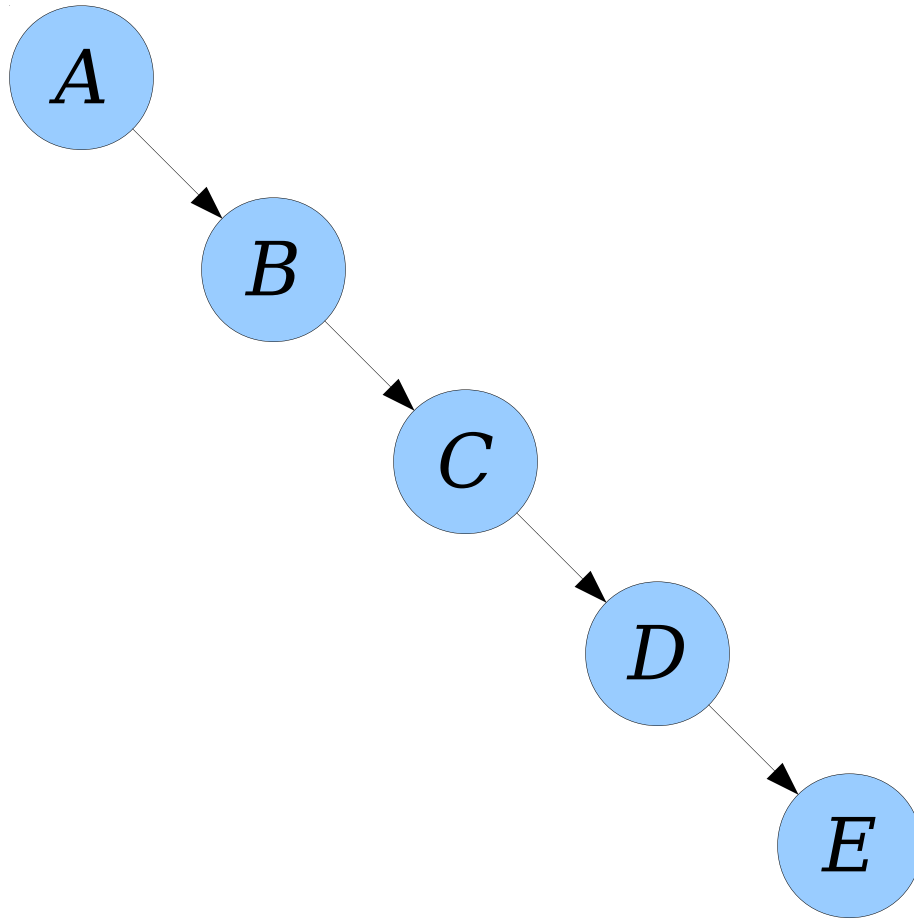

Why Rotate-to-Root Fails

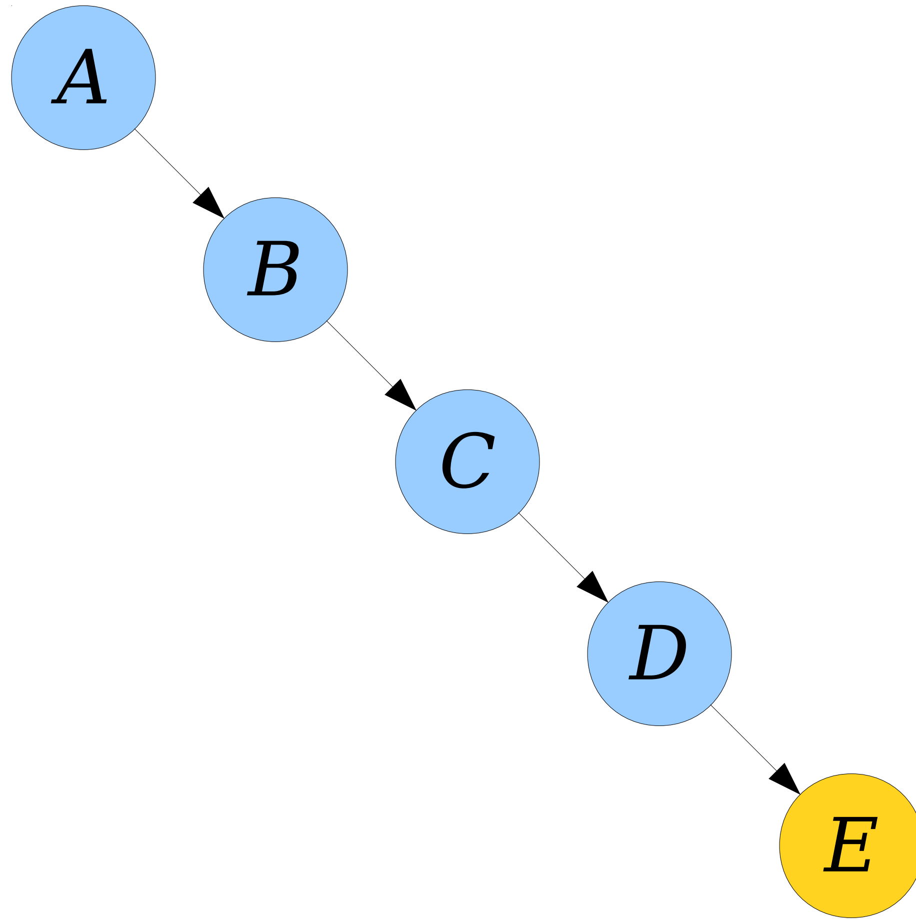

Why Rotate-to-Root Fails

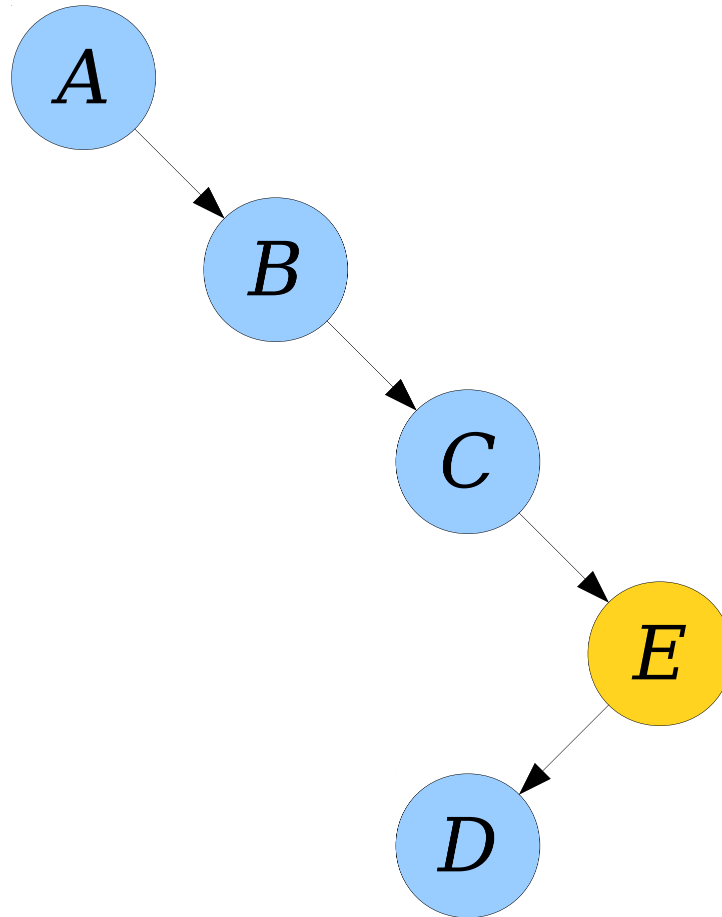

Why Rotate-to-Root Fails

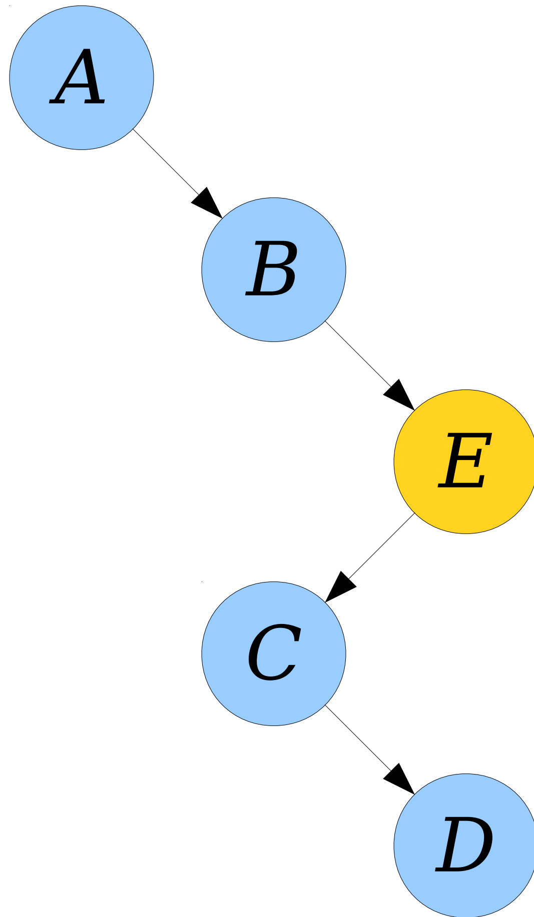

Why Rotate-to-Root Fails

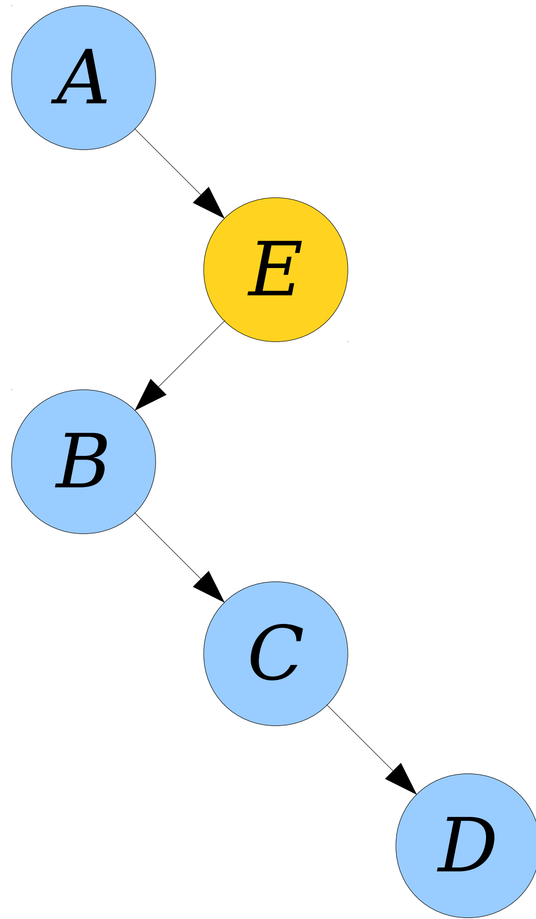

Why Rotate-to-Root Fails

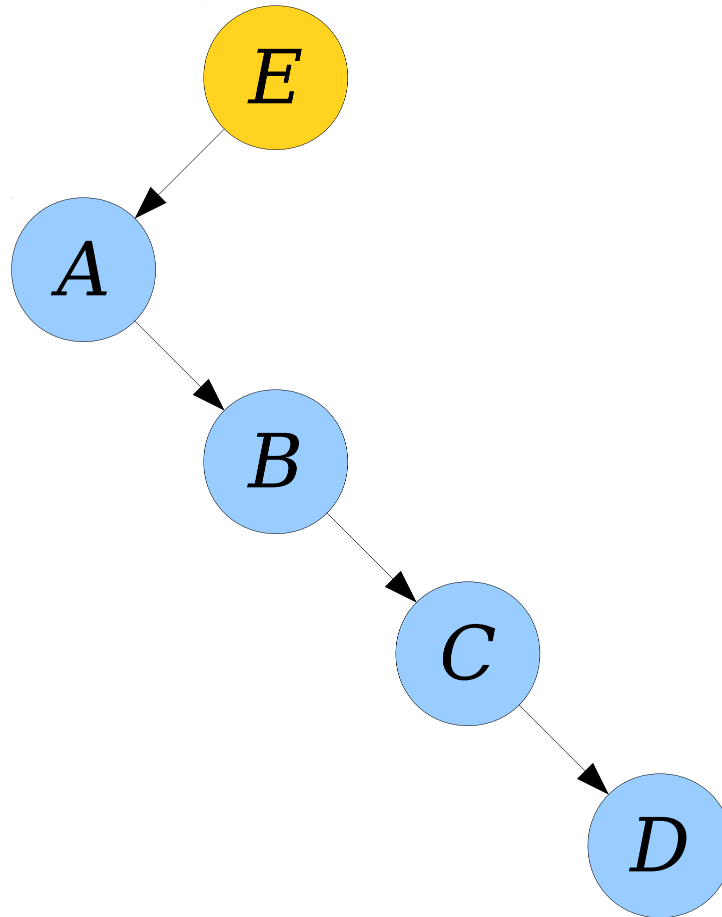

Why Rotate-to-Root Fails

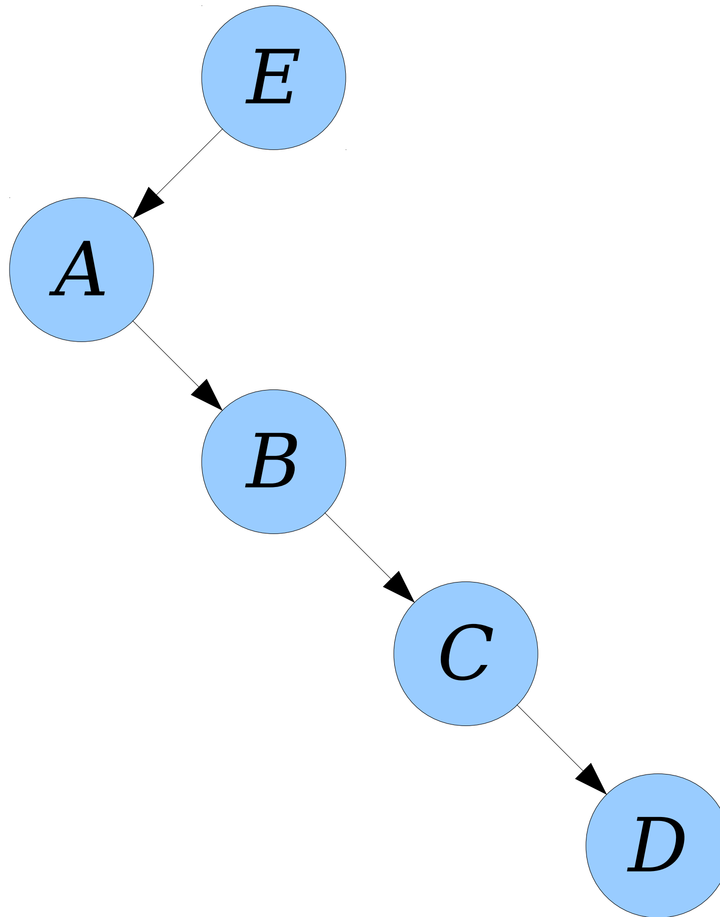

Why Rotate-to-Root Fails

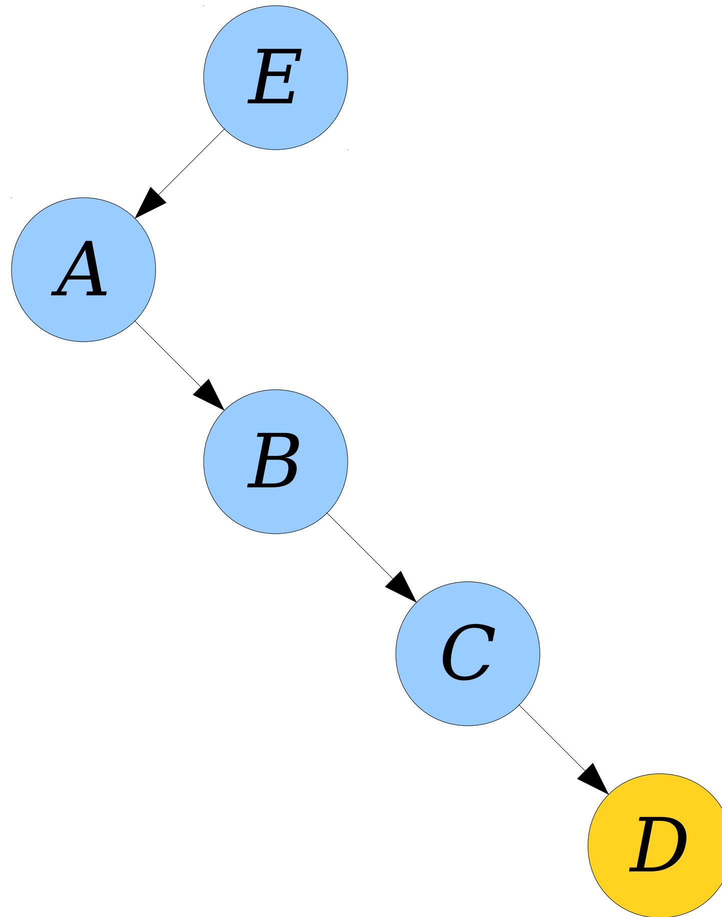

Why Rotate-to-Root Fails

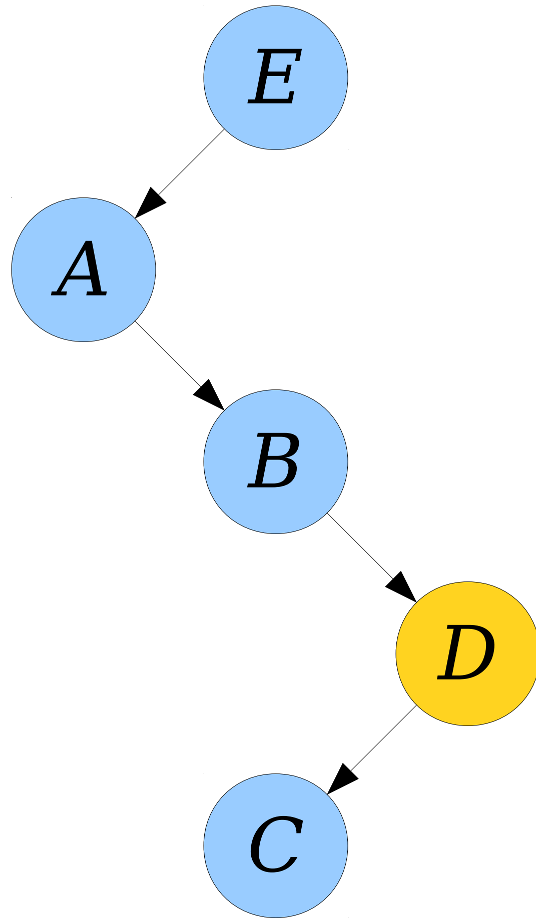

Why Rotate-to-Root Fails

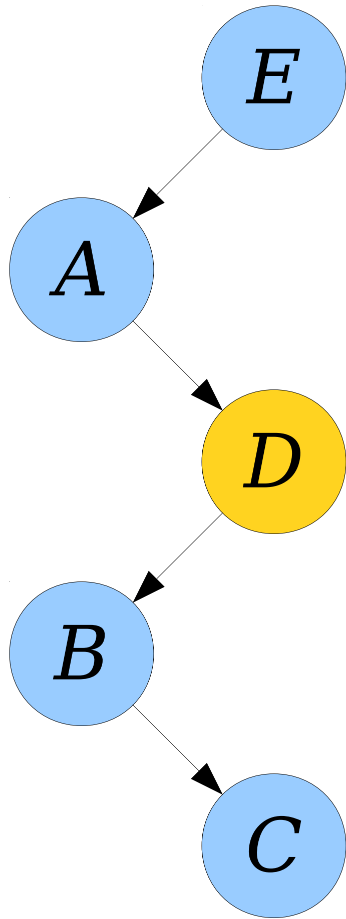

Why Rotate-to-Root Fails

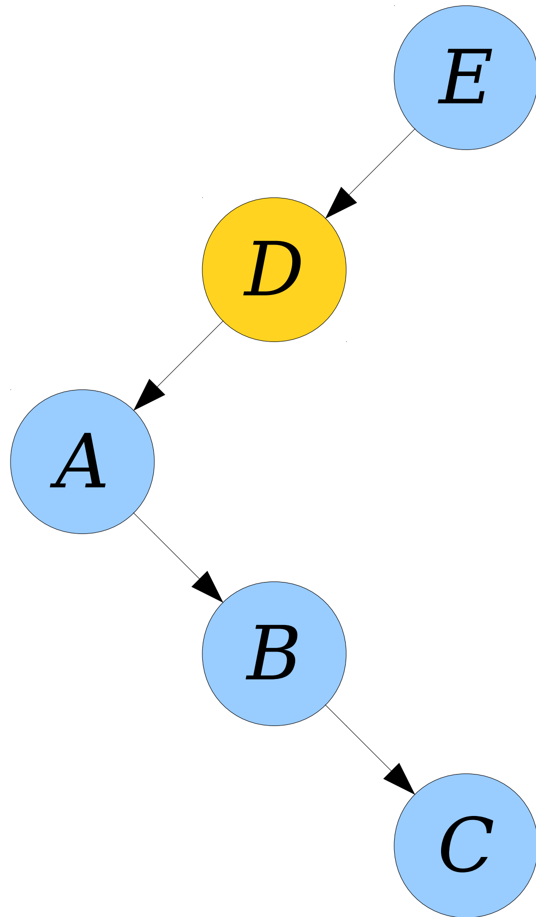

Why Rotate-to-Root Fails

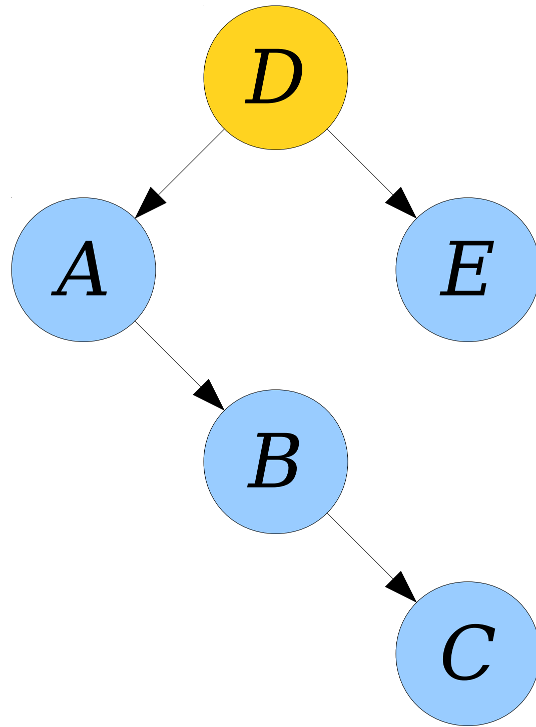

Why Rotate-to-Root Fails

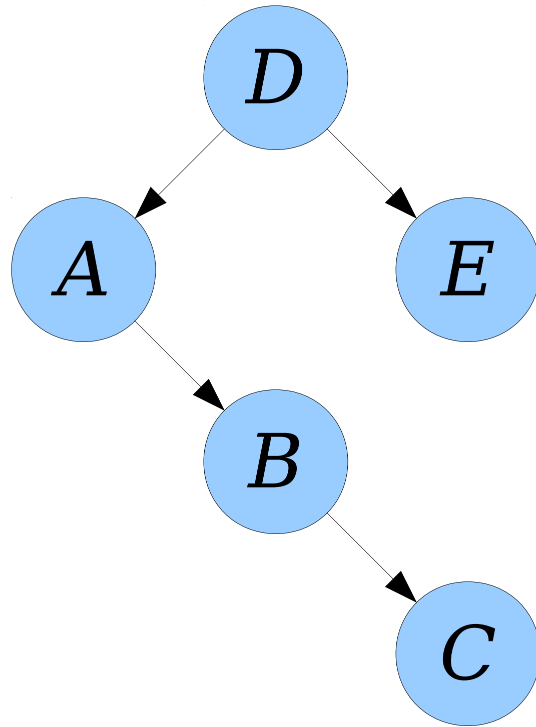

Why Rotate-to-Root Fails

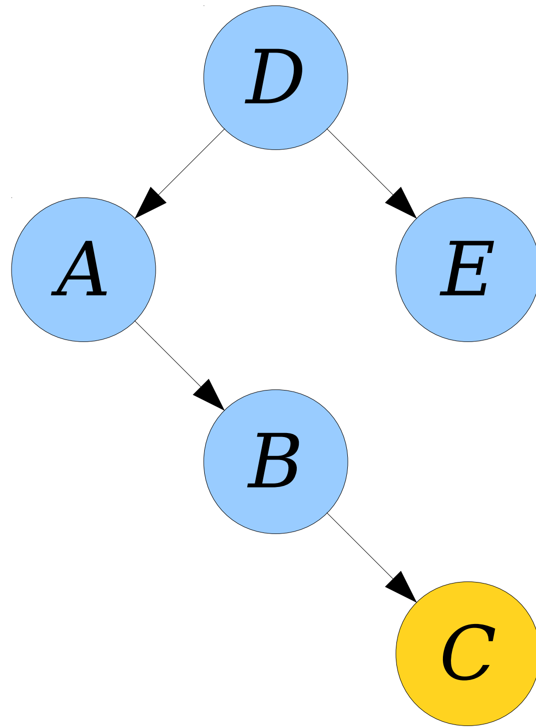

Why Rotate-to-Root Fails

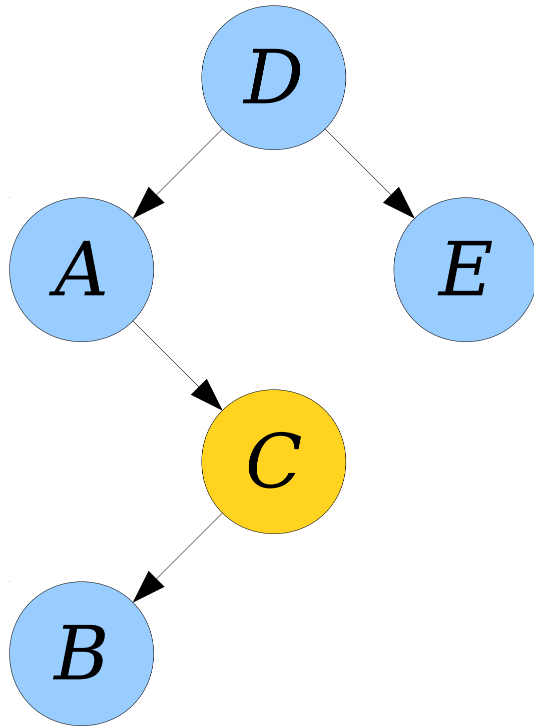

Why Rotate-to-Root Fails

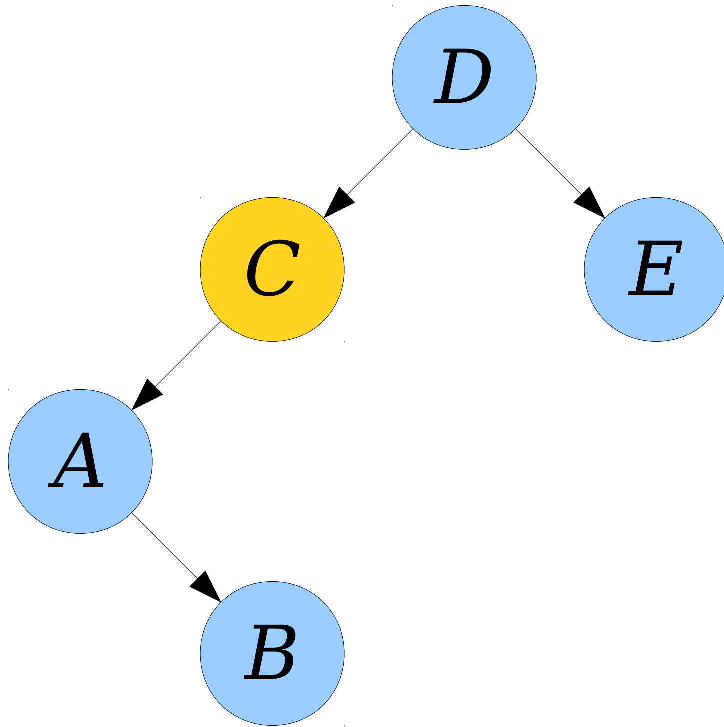

Why Rotate-to-Root Fails

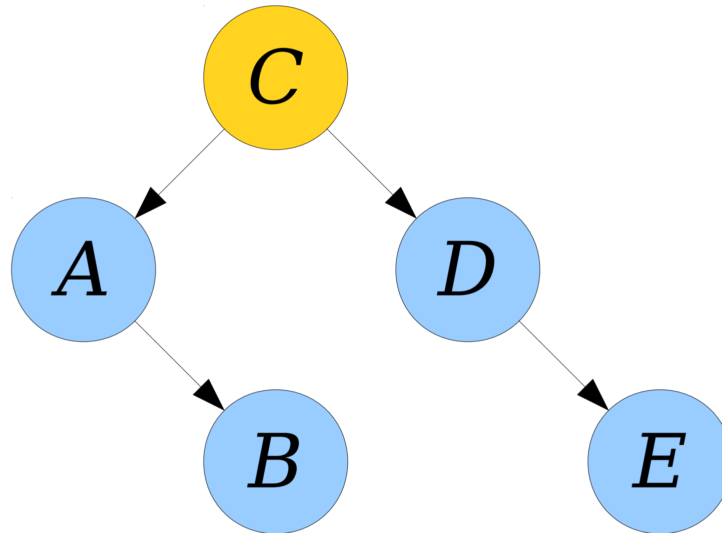

Why Rotate-to-Root Fails

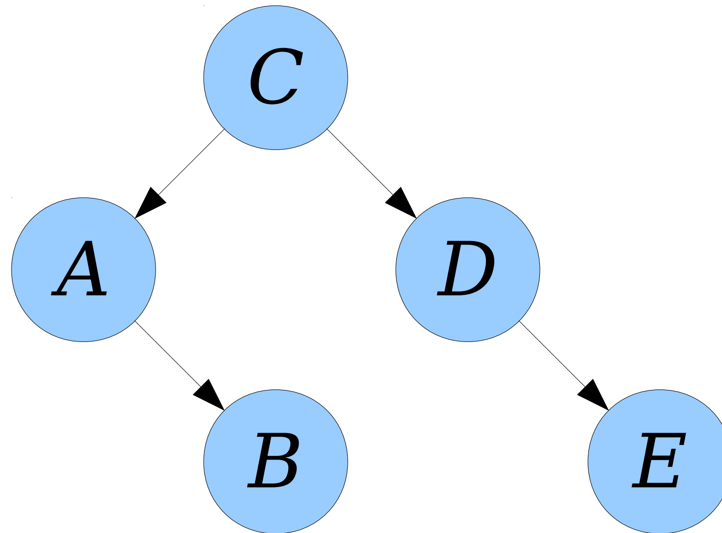

The Wonderful World of Splaying

Some Claims

- ***Claim 1:*** The amortized cost of splaying a node up to the root is O(log n).
- ***Claim 2:*** The amortized cost of splaying a node up to the root can be o(log n) if the access pattern is non-uniform.
- We'll prove these results later today.

Making Things Easy

- Splay trees provide make it extremely easy to perform the following operations:
  - lookup
  - insert
  - delete
  - predecessor / successor
  - join
  - split
- Let's see why.

Lookups

- To do a lookup in a splay tree:
  - Search for that item as usual.
  - If it's found, splay it up to the root.
  - Otherwise, splay the last-visited node to the root.

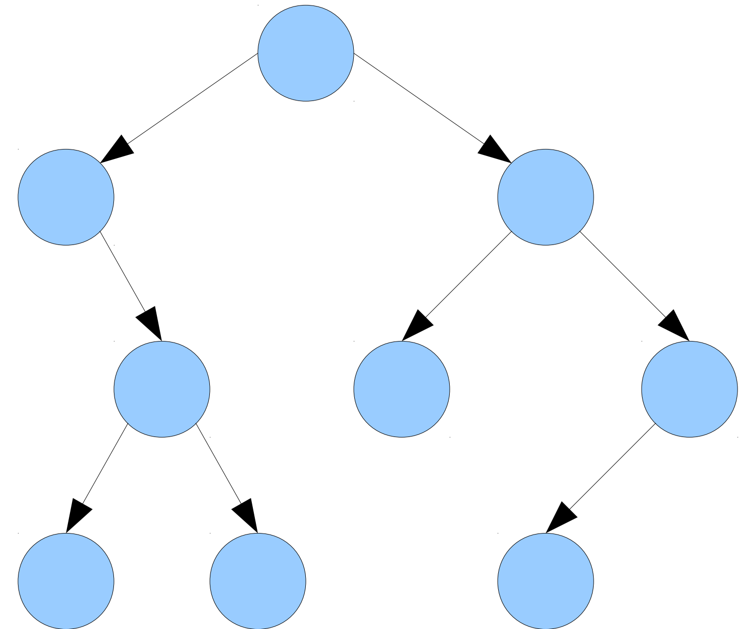

Lookups

- To do a lookup in a splay tree:
  - Search for that item as usual.
  - If it's found, splay it up to the root.
  - Otherwise, splay the last-visited node to the root.

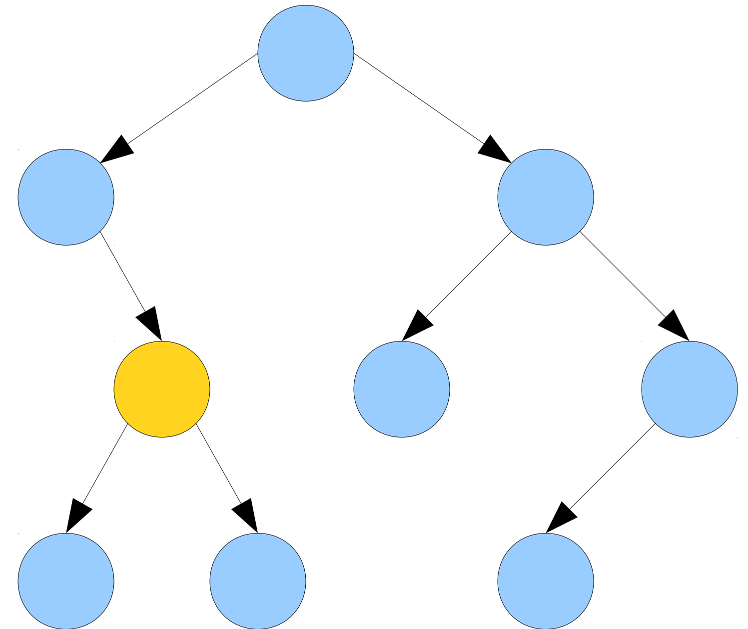

Lookups

- To do a lookup in a splay tree:
  - Search for that item as usual.
  - If it's found, splay it up to the root.
  - Otherwise, splay the last-visited node to the root.

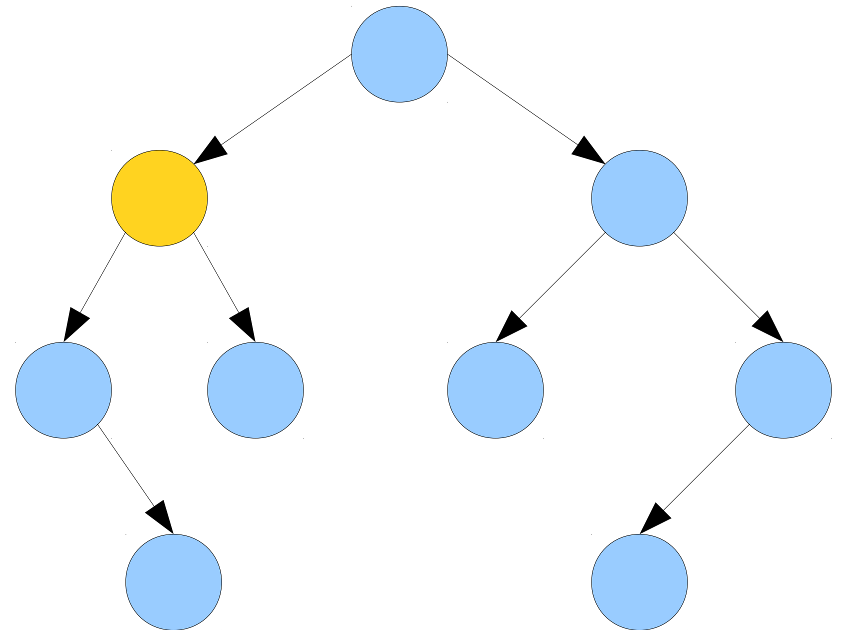

Lookups

- To do a lookup in a splay tree:
  - Search for that item as usual.
  - If it's found, splay it up to the root.
  - Otherwise, splay the last-visited node to the root.

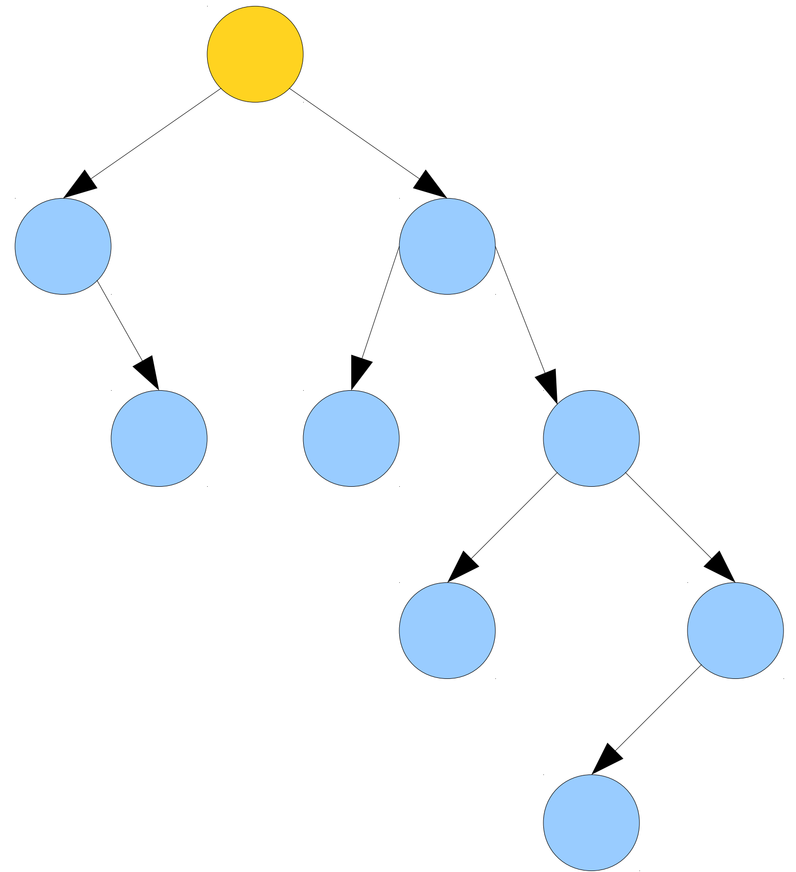

Lookups

- To do a lookup in a splay tree:
  - Search for that item as usual.
  - If it's found, splay it up to the root.
  - Otherwise, splay the last-visited node to the root.

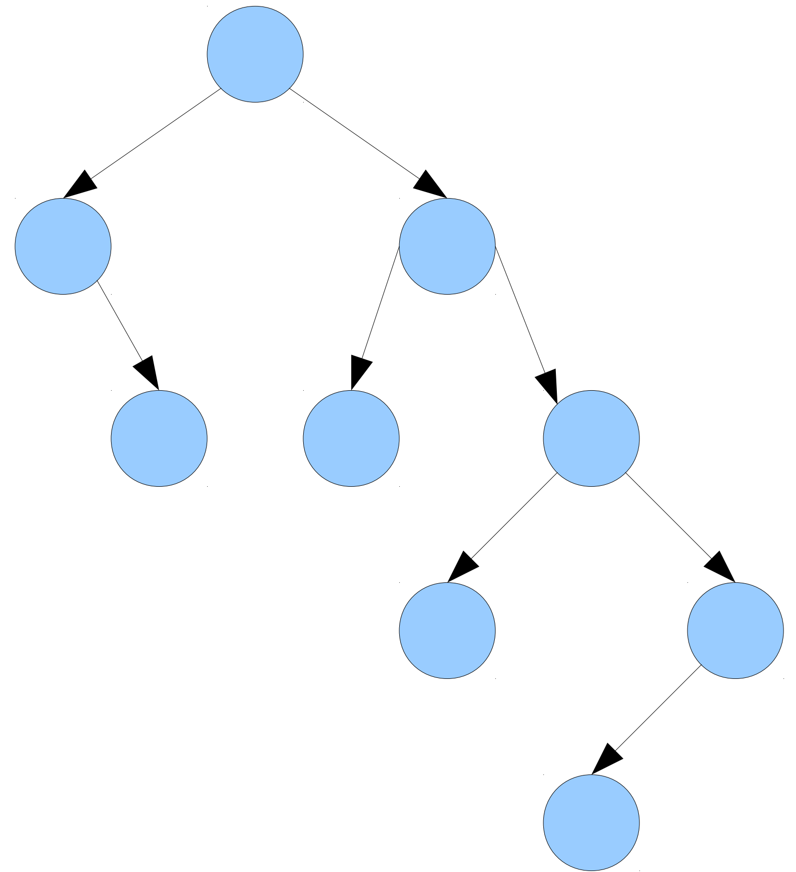

Insertions

- To insert a node into a splay tree:
  - Insert the node as usual.
  - Splay it up to the root.

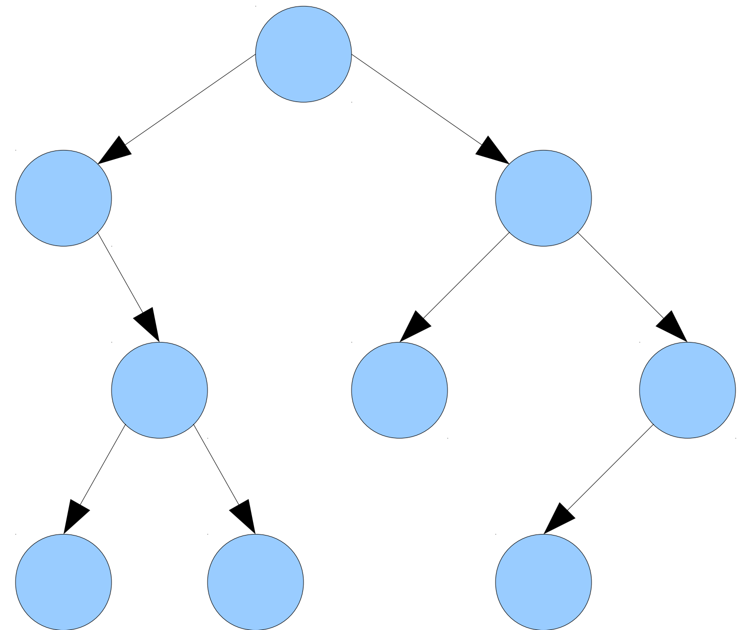

Insertions

- To insert a node into a splay tree:
  - Insert the node as usual.
  - Splay it up to the root.

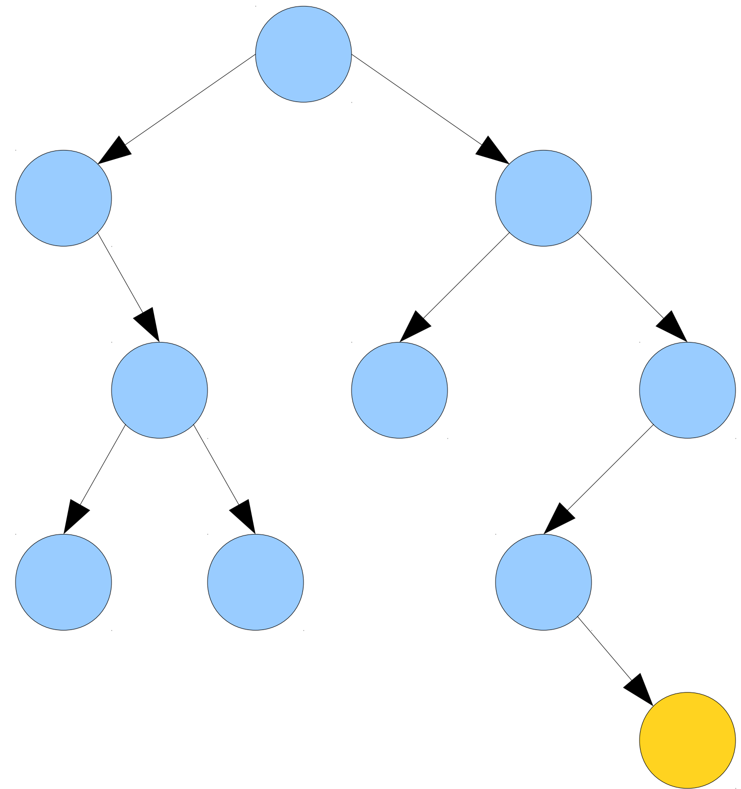

Insertions

- To insert a node into a splay tree:
  - Insert the node as usual.
  - Splay it up to the root.

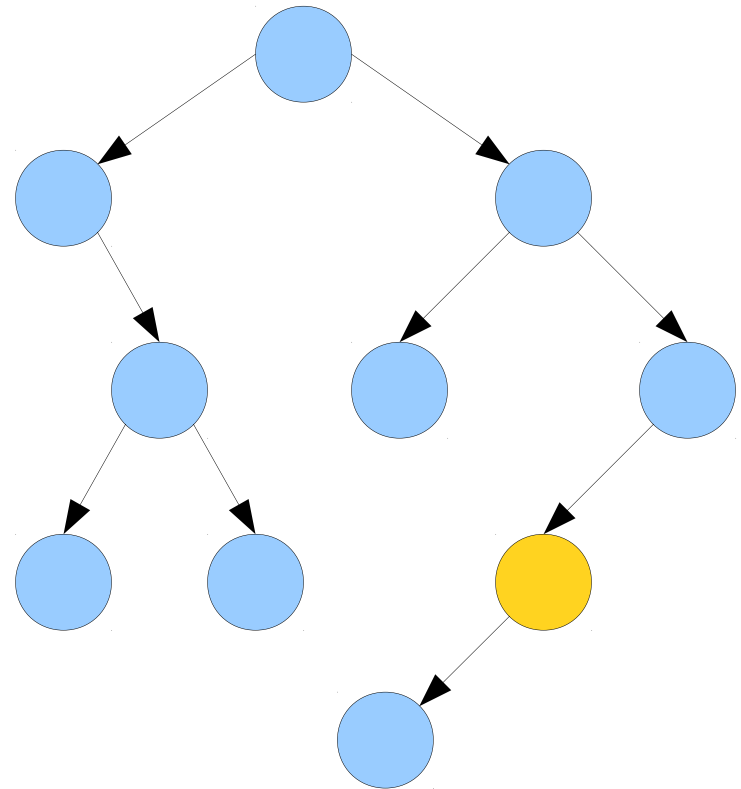

Insertions

- To insert a node into a splay tree:
  - Insert the node as usual.
  - Splay it up to the root.

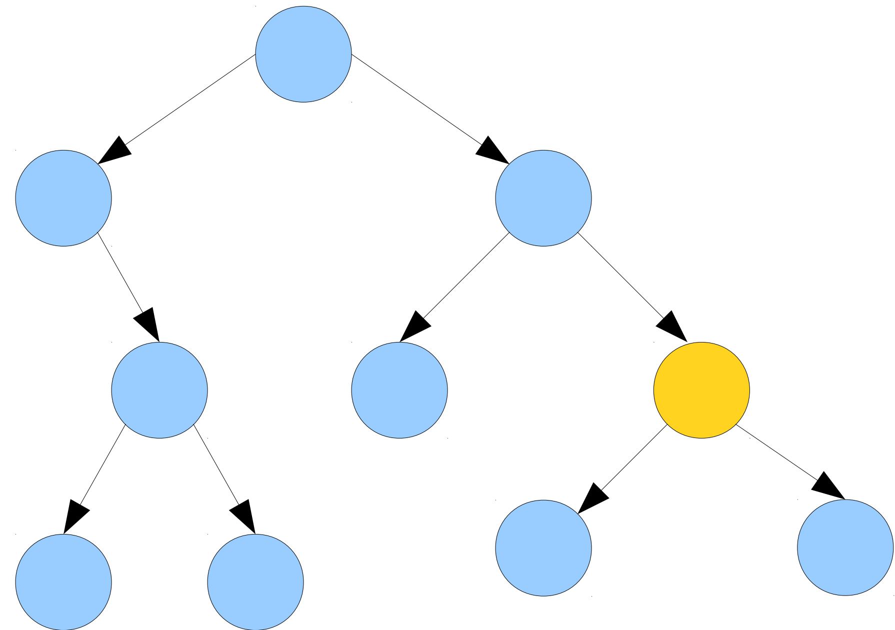

Insertions

- To insert a node into a splay tree:
  - Insert the node as usual.
  - Splay it up to the root.

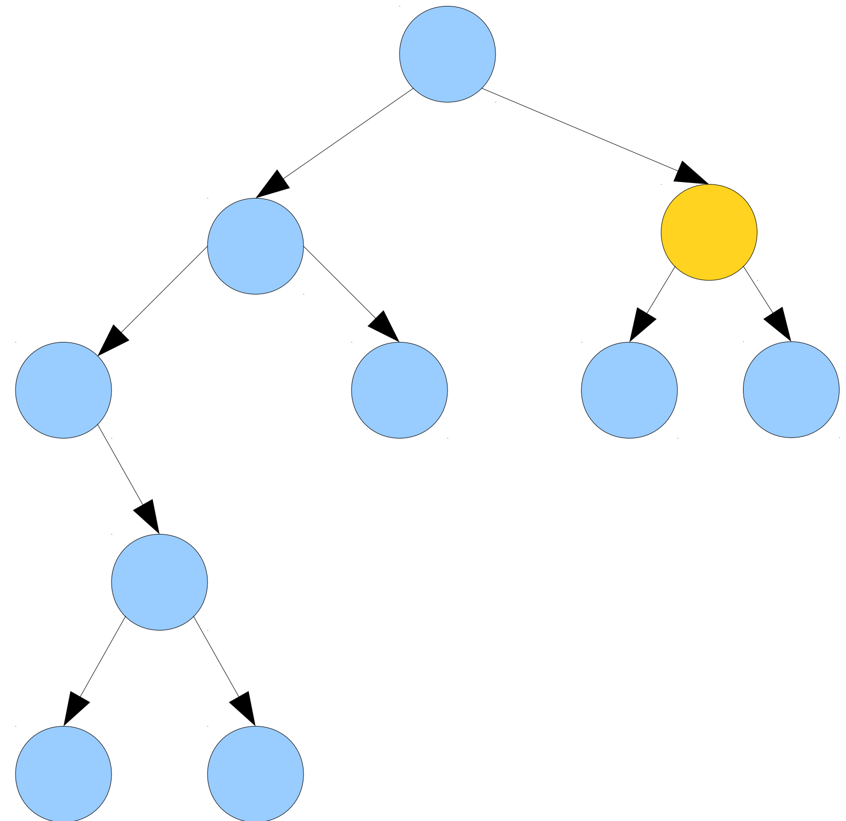

Insertions

- To insert a node into a splay tree:
  - Insert the node as usual.
  - Splay it up to the root.

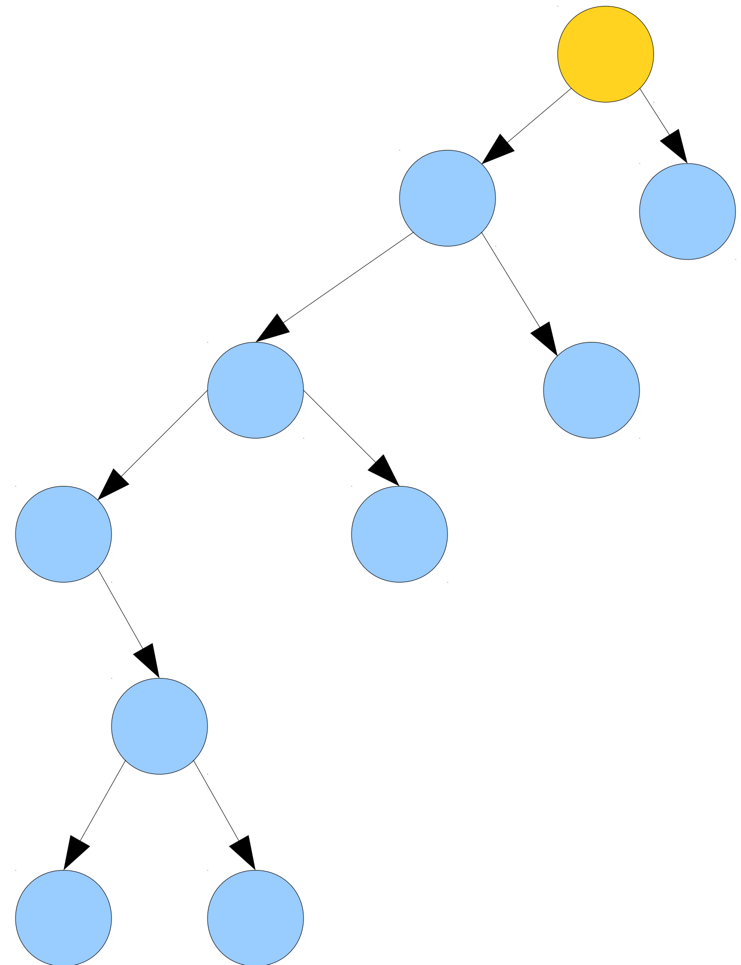

Insertions

- To insert a node into a splay tree:
  - Insert the node as usual.
  - Splay it up to the root.

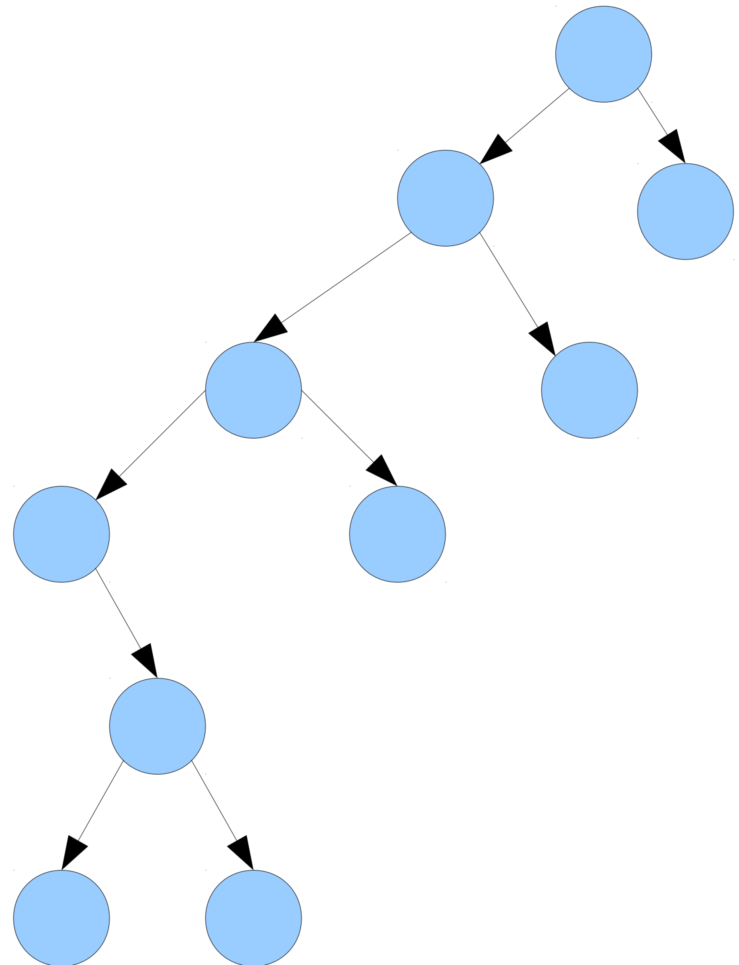

Join

- To join two trees T_1 and T_2, where all keys in T_1 are less than the keys in T_2:
  - Splay the max element of T_1 to the root.
  - Make T_2 a right child of T_1.

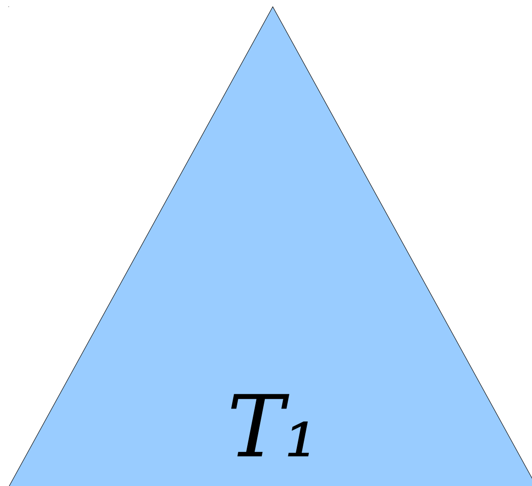

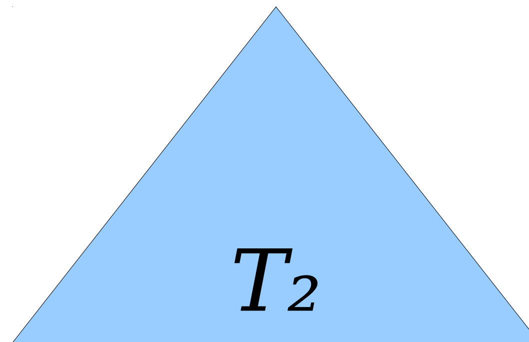

Join

- To join two trees T_1 and T_2, where all keys in T_1 are less than the keys in T_2:
  - Splay the max element of T_1 to the root.
  - Make T_2 a right child of T_1.

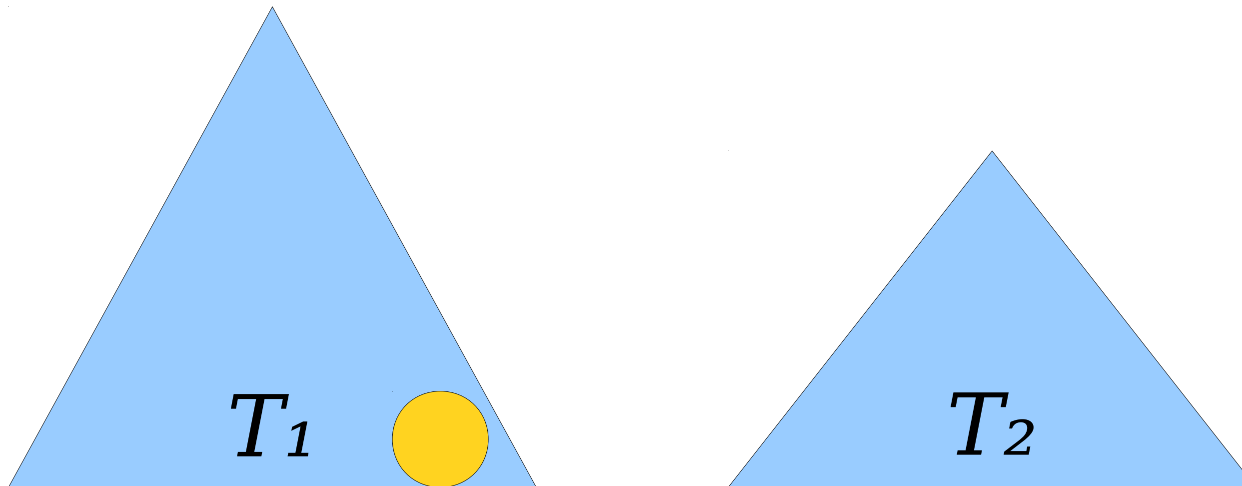

Join

- To join two trees T_1 and T_2, where all keys in T_1 are less than the keys in T_2:
  - Splay the max element of T_1 to the root.
  - Make T_2 a right child of T_1.

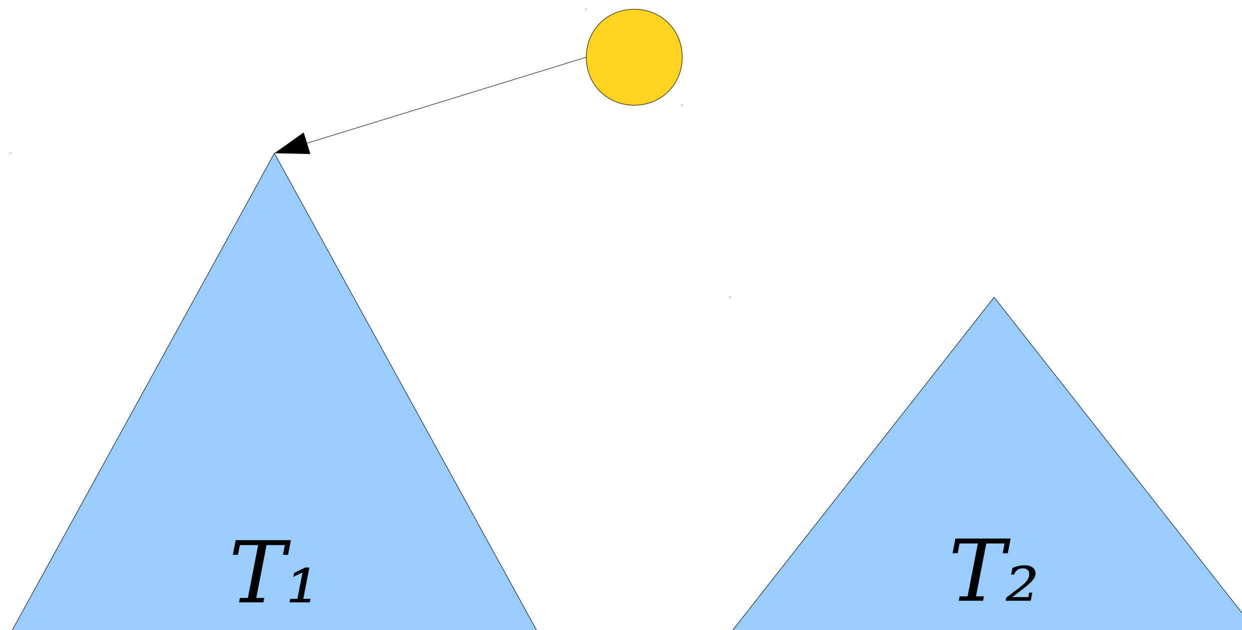

Join

- To join two trees T_1 and T_2, where all keys in T_1 are less than the keys in T_2:
  - Splay the max element of T_1 to the root.
  - Make T_2 a right child of T_1.

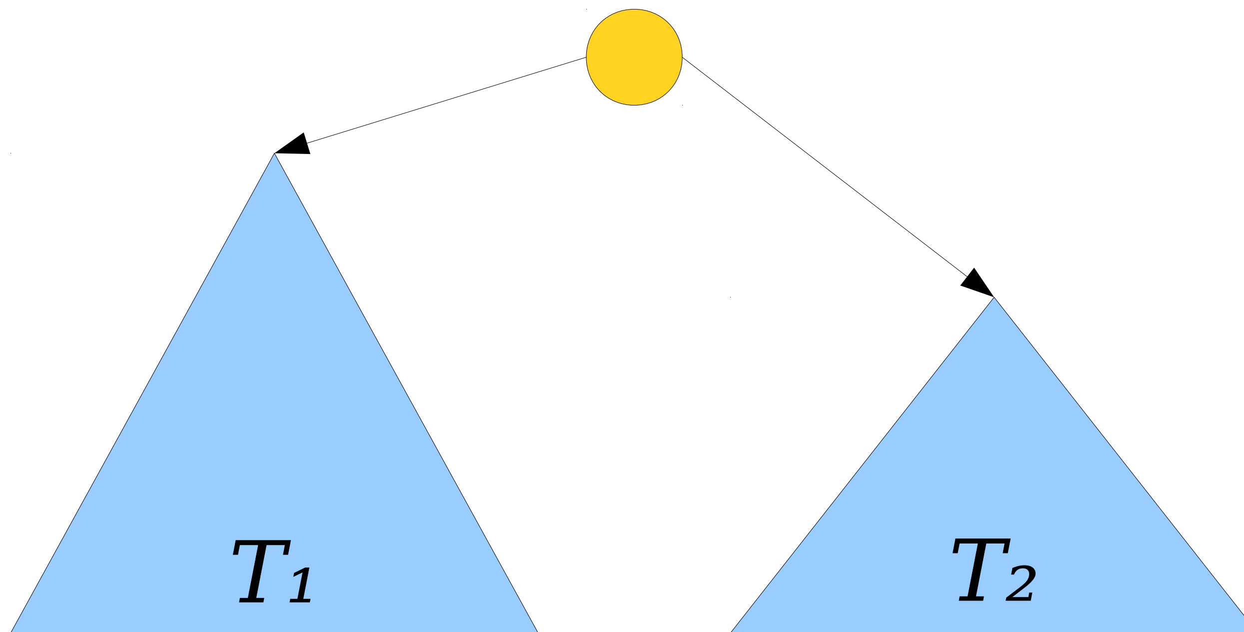

Split

- To split T at a key k:
  - Splay the successor of k up to the root.
  - Cut the link from the root to its left child.

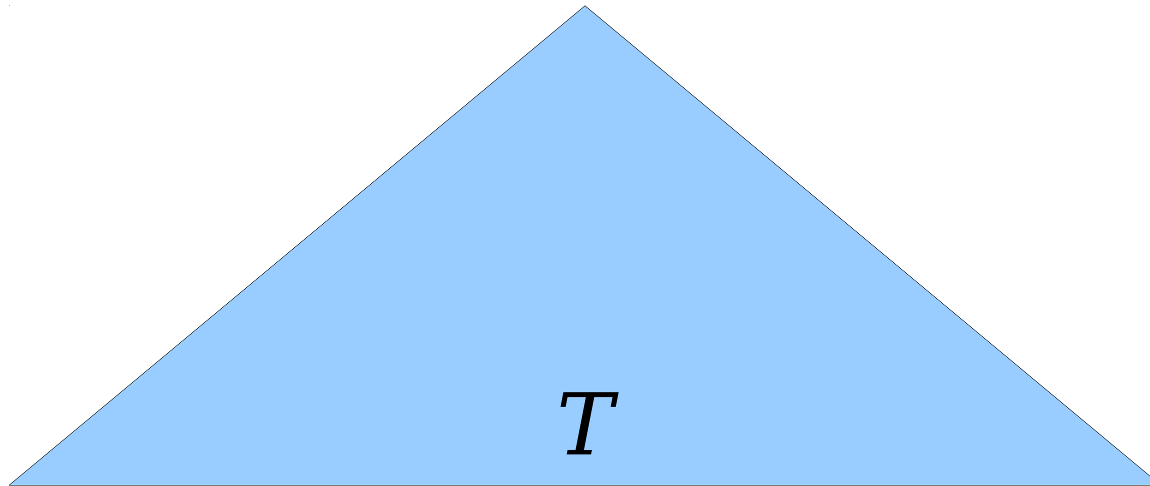

Split

- To split T at a key k:
  - Splay the successor of k up to the root.
  - Cut the link from the root to its left child.

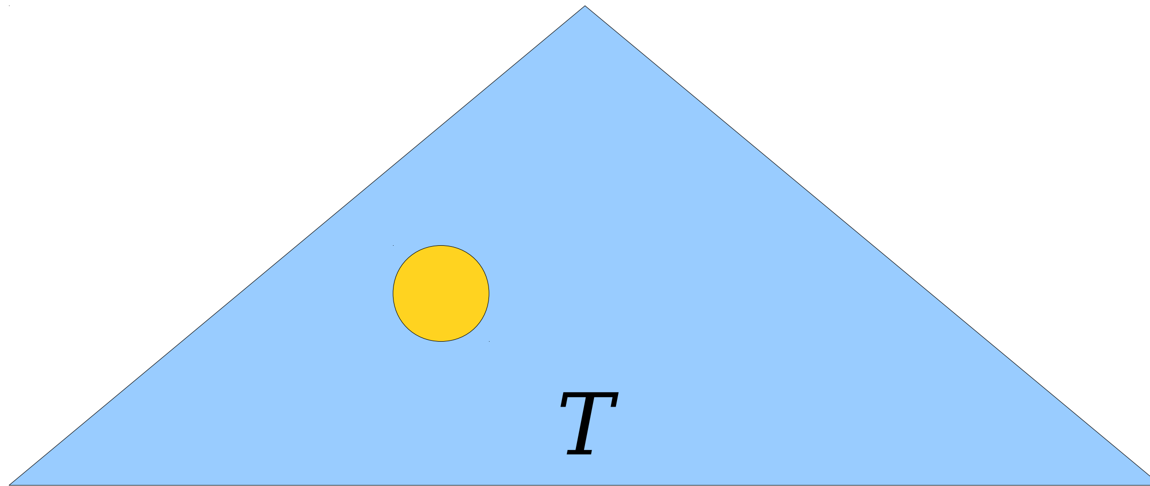

Split

- To split T at a key k:
  - Splay the successor of k up to the root.
  - Cut the link from the root to its left child.

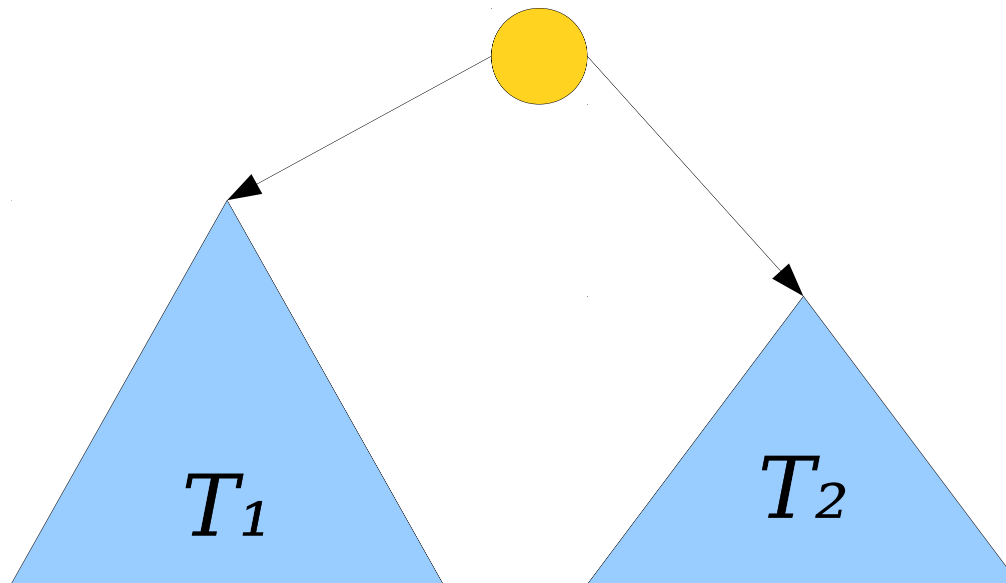

Split

- To split T at a key k:
  - Splay the successor of k up to the root.
  - Cut the link from the root to its left child.

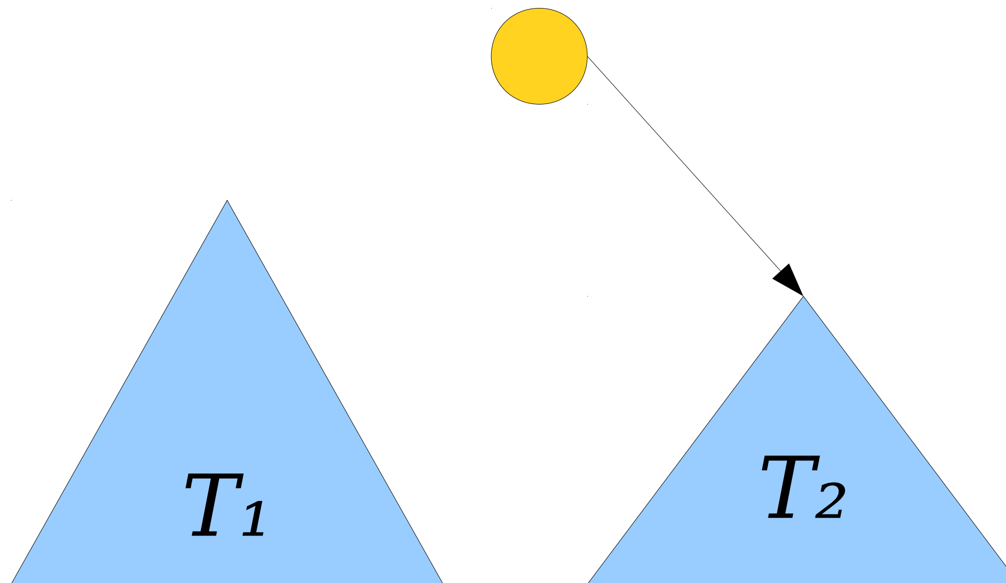

Delete

- To delete a key k from the tree:
  - Splay k to the root.
  - Delete k.
  - Join the two resulting subtrees.

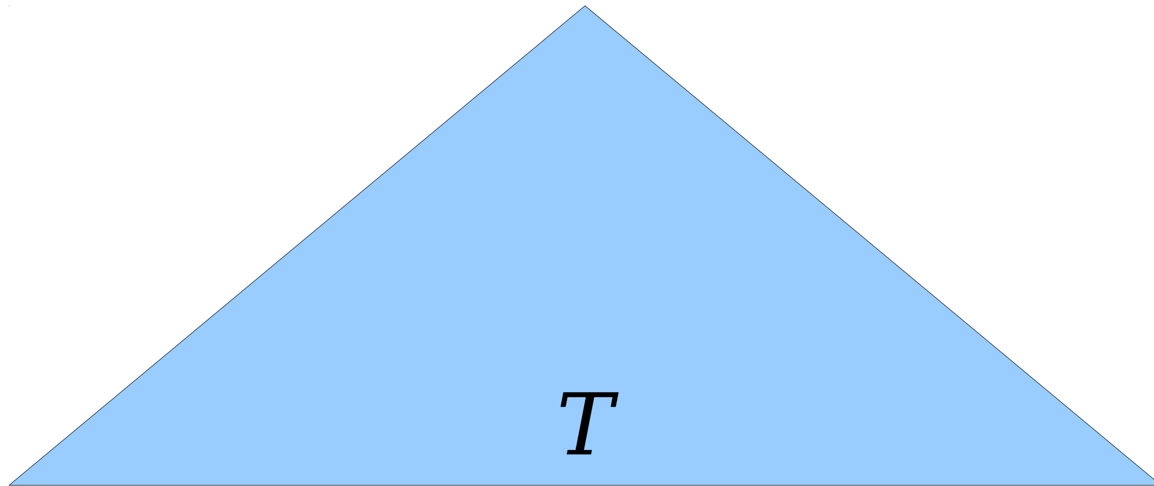

Delete

- To delete a key k from the tree:
  - Splay k to the root.
  - Delete k.
  - Join the two resulting subtrees.

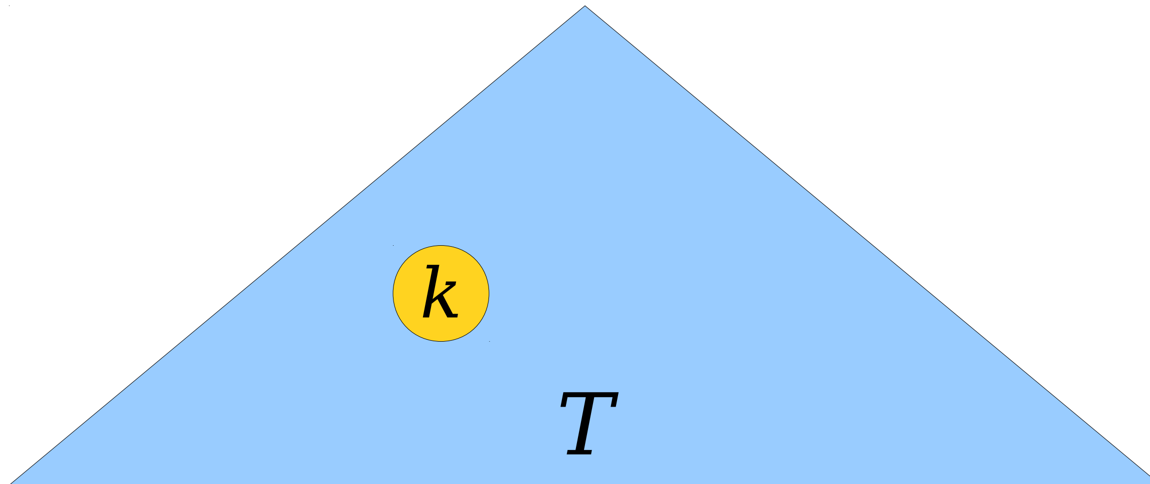

Delete

- To delete a key k from the tree:
  - Splay k to the root.
  - Delete k.
  - Join the two resulting subtrees.

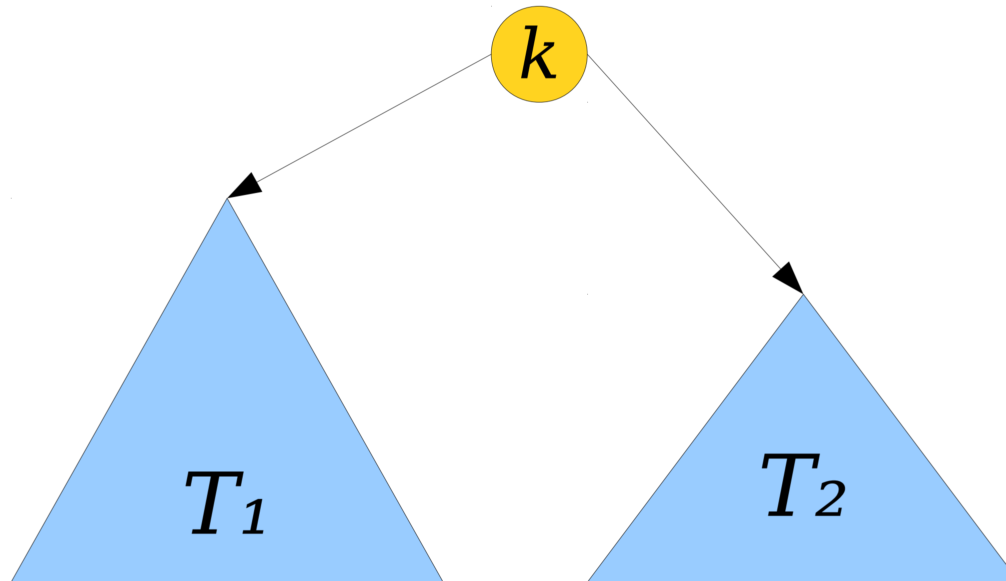

Delete

- To delete a key k from the tree:
  - Splay k to the root.
  - Delete k.
  - Join the two resulting subtrees.

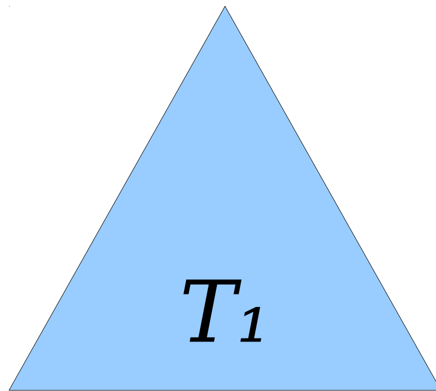

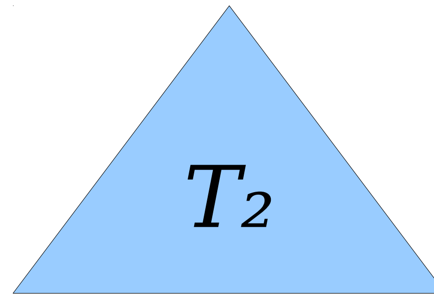

Delete

- To delete a key k from the tree:
  - Splay k to the root.
  - Delete k.
  - Join the two resulting subtrees.

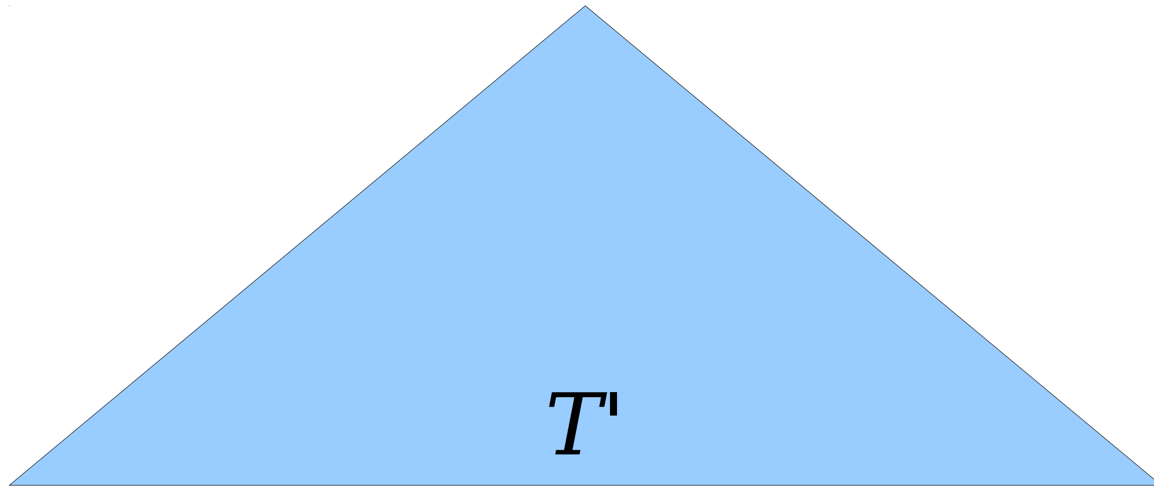

The Runtime

- ***Claim:*** All of these operations require amortized time O(log n).
- ***Rationale:*** Each has runtime bounded by the cost of O(1) splays, which takes total amortized time O(log n).
- Contrast this with red/black trees:
  - No need to store any kind of balance information.
  - Only three rules to memorize.

So… just how fast *are* splay trees?

Time-Out for Announcements!

The Stanford Women in Computer Science EXEC Application is CURRENTLY LIVE

Interested in being an influential promoter of women in computer science?

Want to shape the future of Stanford's computer science community?

How about meeting some of the most talented CS students at Stanford?

You've come to the right place! ALL genders welcome to apply!

Check out our chair positions for next year's program and make sure to apply by Friday, May 11 at 11:59pm!

If you have questions, please don't hesitate to reach out to Lauren Zhu (laurenz@stanford.edu) and Makena Low (makenal@stanford.edu).

Problem Set Four

- Problem Set Four goes out today. It's due next Thursday at 2:30PM.
  - Play around with amortized analyses, binomial heaps, Fibonacci heaps, and splay trees!
- Problem Set Three solutions are now up on the course website.
- Have questions? Feel free to stop by office hours or to ask on Piazza!

Final Project Logistics

- As a reminder, the final project proposal is due this Thursday at 2:30PM.
- ***No late days may be used here***. We'd like to get our matchmaking done as early as possible.
- There's a huge list of potential topics up on the course website. Feel free to read over it for inspiration!

Back to CS166!

The Tricky Part: Formalizing This

Analyzing BSTs

- Let's assume that every element x_i in our BST has some associated weight w_i.
- We'll assume that access probabilities are proportional to weights, though we won't assume the weights sum to 1. (This is just for mathematical simplicity.)
- Let W denote the sum $w_1 + w_2 + \ldots + w_n$.
- Imagine we have some fixed BST T containing the keys $x_1, \ldots, x_n$. Let s_i denote the sum of all the weights of the keys in the subtree rooted at x_i.

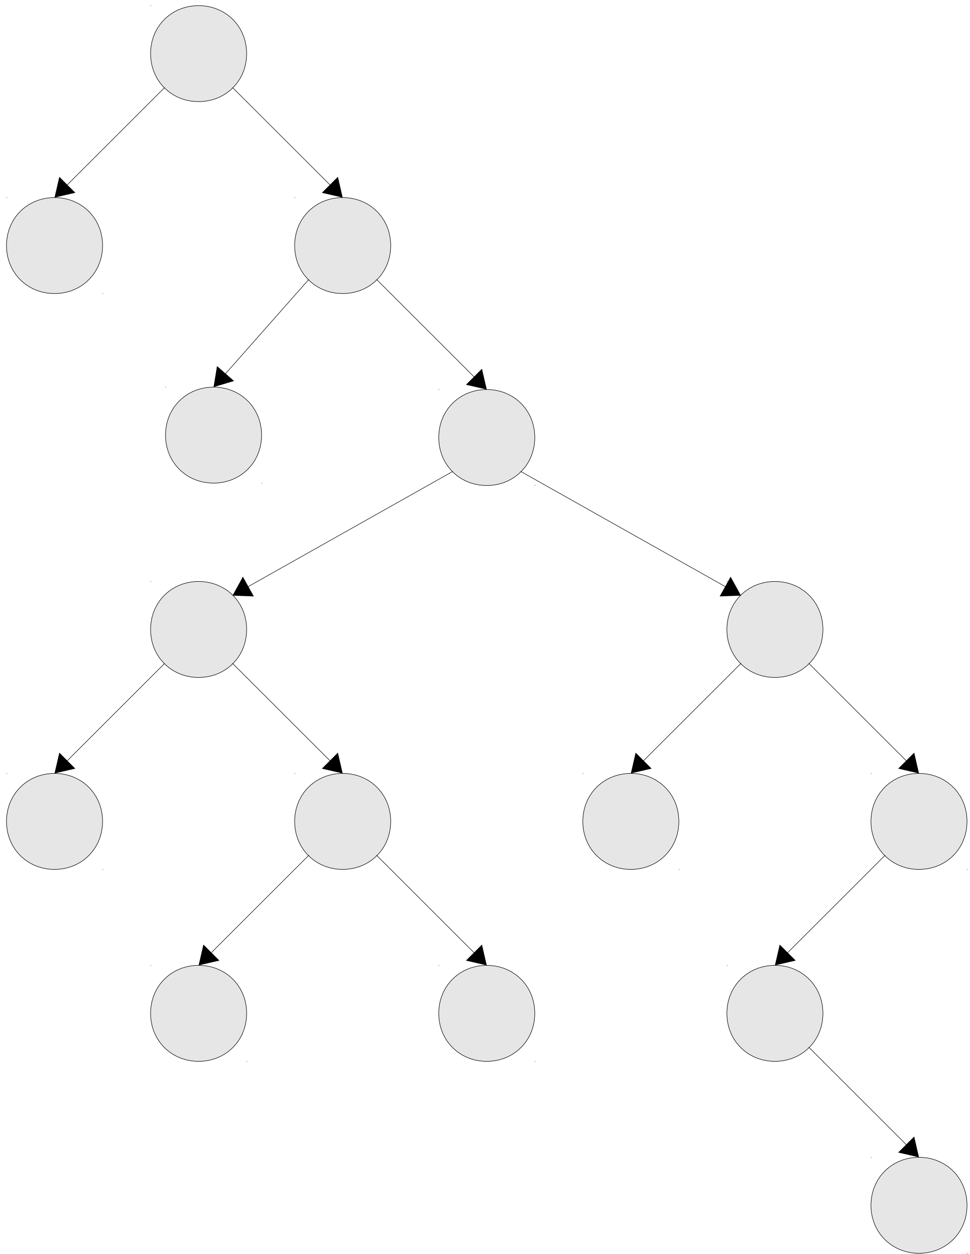

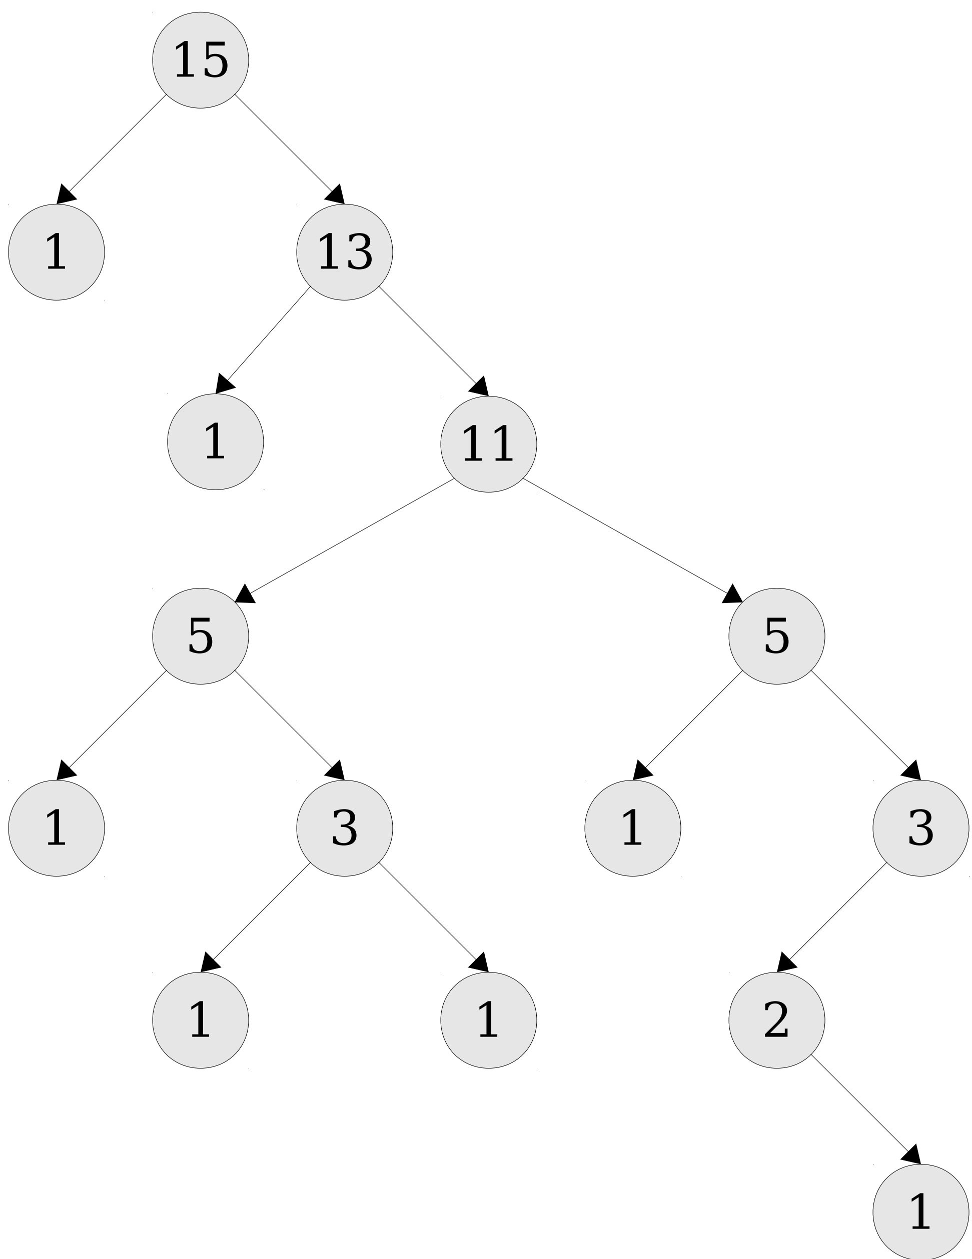
15
1
13
1
11
5
5
1
3
1
3
1
1
2
1

$s_{new} > \frac{1}{2}s_{old}$

$s_{new} \leq \frac{1}{2}s_{old}$

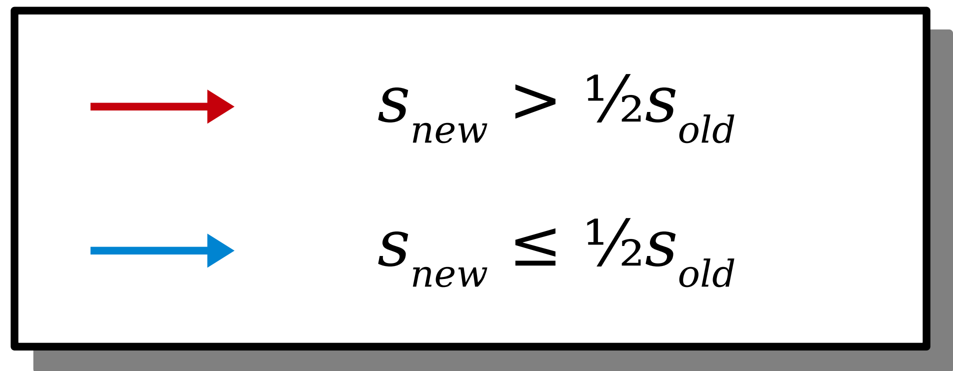
$s_{new} > \frac{1}{2} s_{old}$
$s_{new} \leq \frac{1}{2} s_{old}$

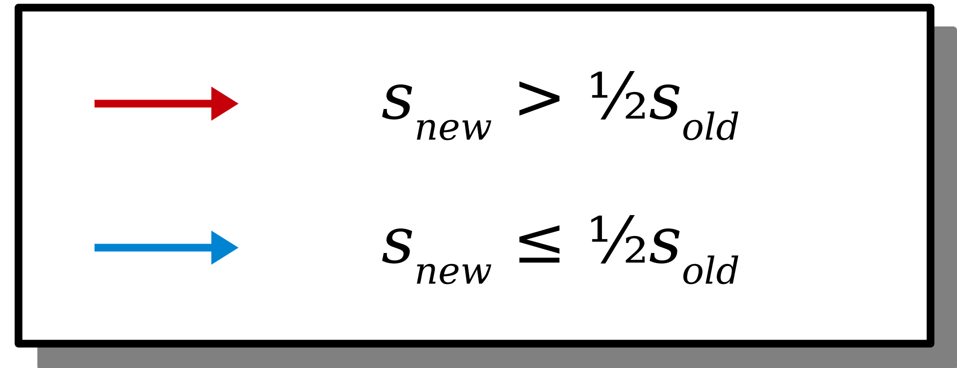

Cost of looking up x_i:

O(**#blue-used** + **#red-used**)

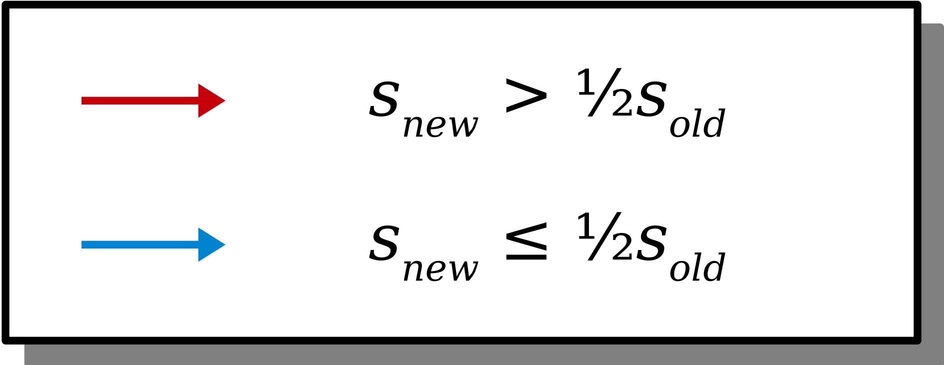

Cost of looking up x_i:

O(**log (W / w_i)** + **#red-used**)

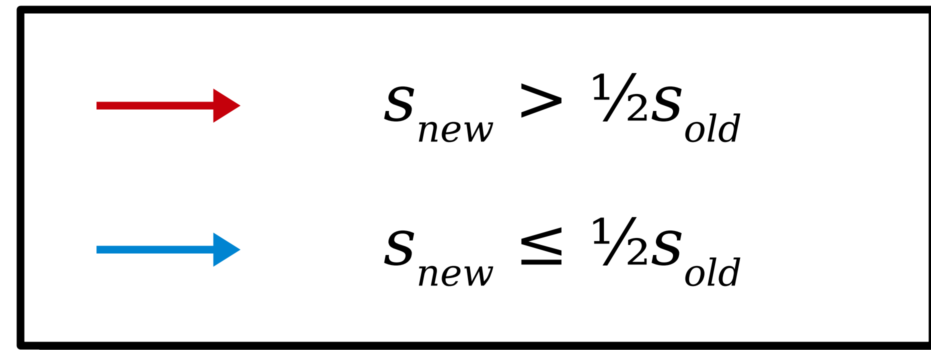

Cost of looking up x_i:

O(**log (*W* / *w*$_i$)** + **#red-used**)

Every blue edge throws away half of the total weight remaining.

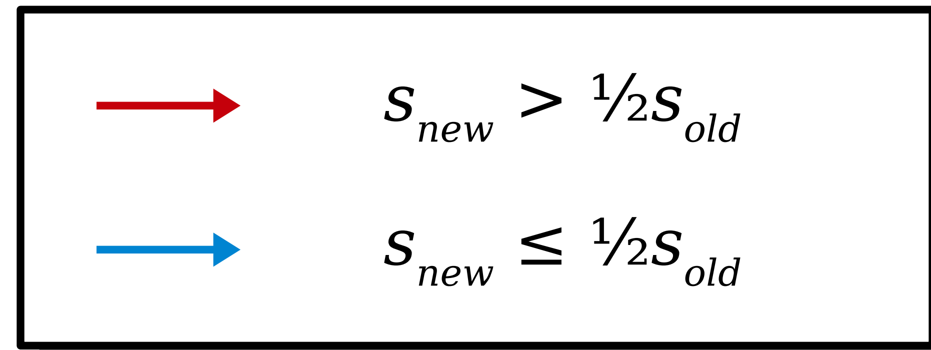

Cost of looking up x_i:

$O($**log (*W* / *w*$_i$)** + **#red-used**$)$

Every blue edge throws away half of the total weight remaining.

We begin with W total weight. If we're searching for x_i, the total weight at node x_i must be at least w_i.

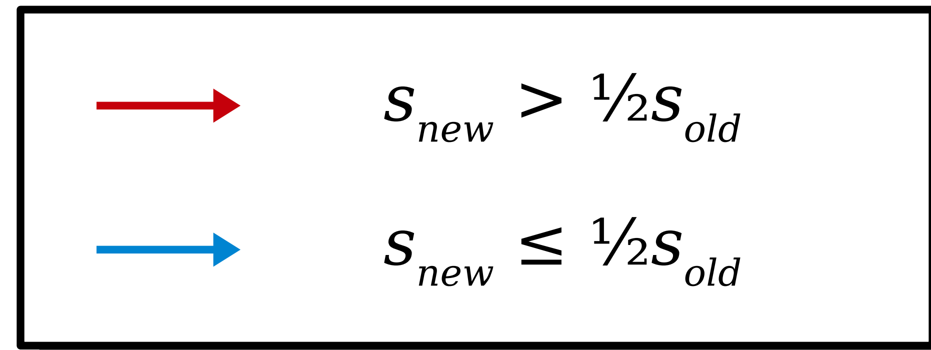

Cost of looking up x_i:

O(**log (W / w_i)** + **#red-used**)

Every blue edge throws away half of the total weight remaining.

We begin with W total weight. If we're searching for x_i, the total weight at node x_i must be at least w_i.

You can only throw away half the total weight log (W / w_i) times before the total weight drops to w_i.

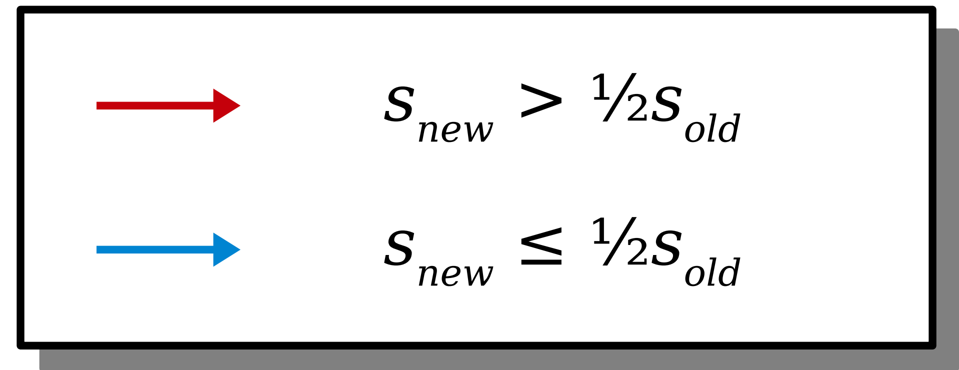

15
1
13
1
11
5
5
1
3
1
3
1
1
2
1

Cost of looking up x_i:

O(**log (*W* / *w*$_i$)** + **#red-used**)

This technique of splitting edges into "good" edges and "bad" edges is a technique called a ***heavy/light decomposition*** and is used extensively in the design and analysis of data structures.

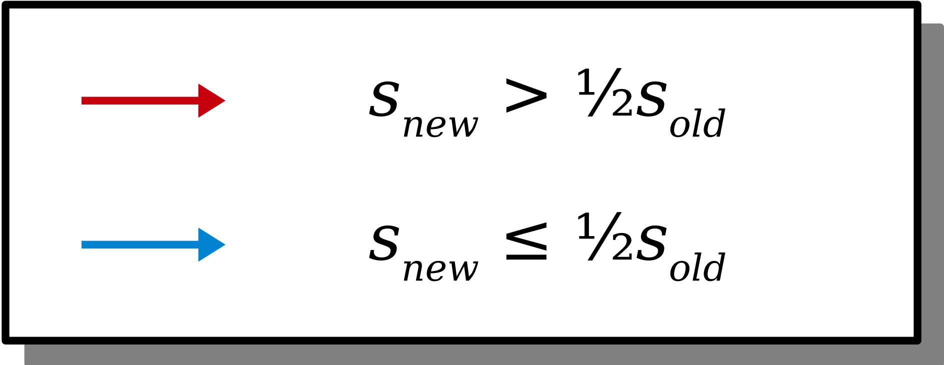

Cost of looking up x_i:

O(**log (W / w_i)** + **#red-used**)

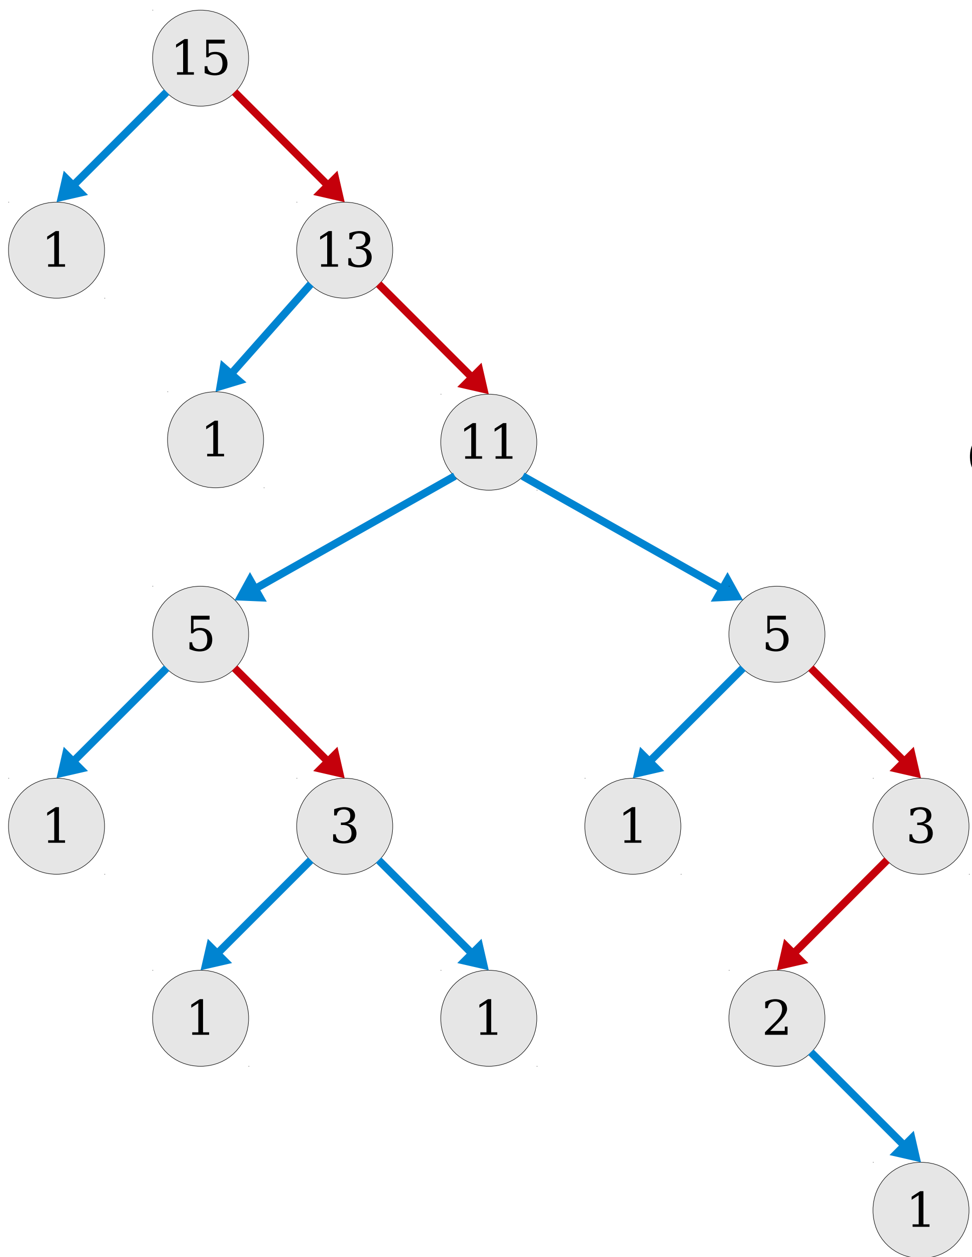

$\lg s_{new} > \lg (\tfrac{1}{2} s_{old})$

$\lg s_{new} \leq \lg (\tfrac{1}{2} s_{old})$

Cost of looking up x_i:

$O(\mathbf{\log (W / w_i)} + \mathbf{\#red\text{-}used})$

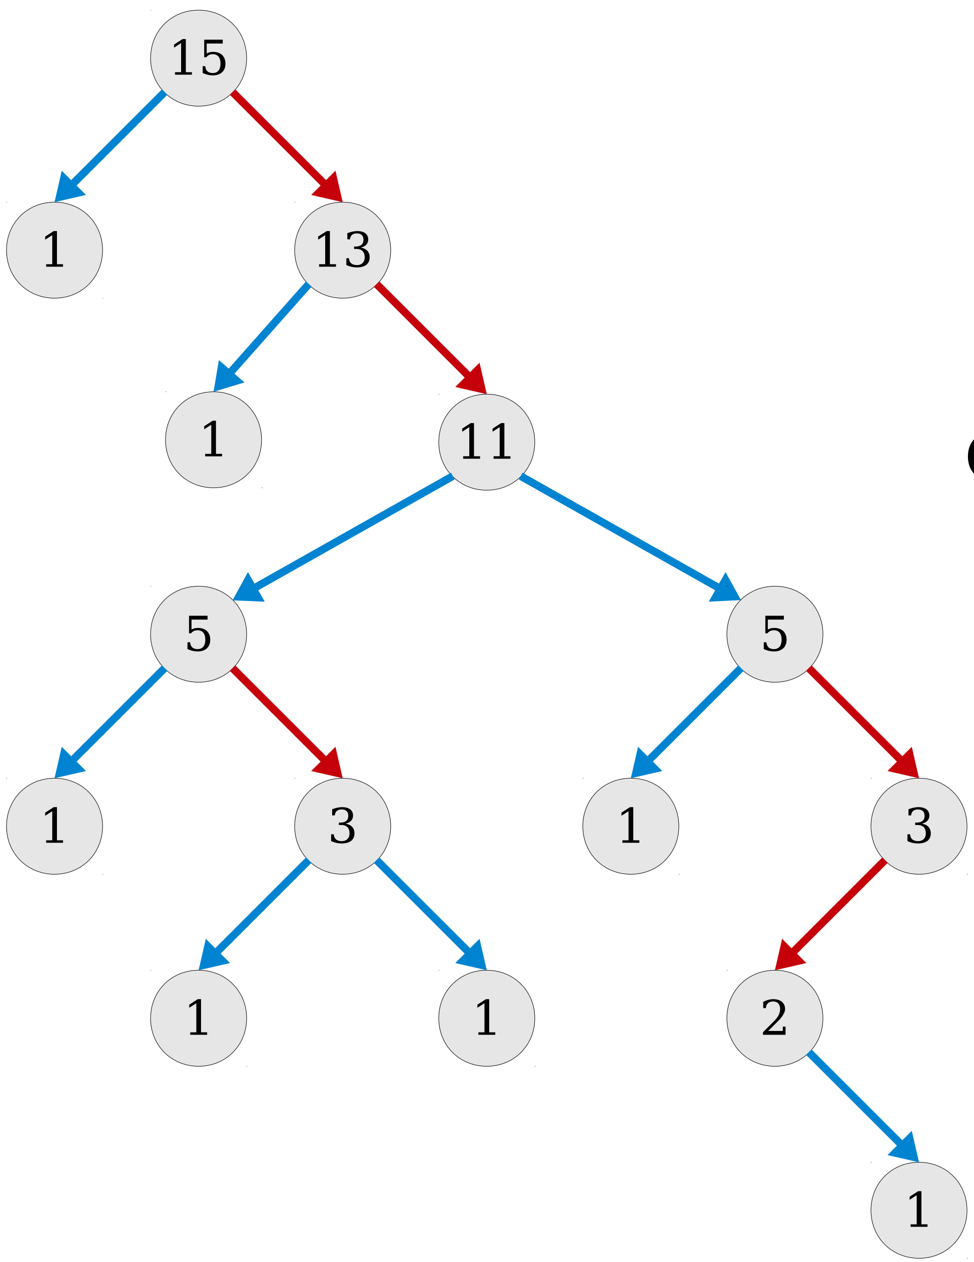

$\lg s_{new} > \lg s_{old} - 1$

$\lg s_{new} \leq \lg s_{old} - 1$

Cost of looking up x_i:

O(**log (W / w_i)** + **#red-used**)

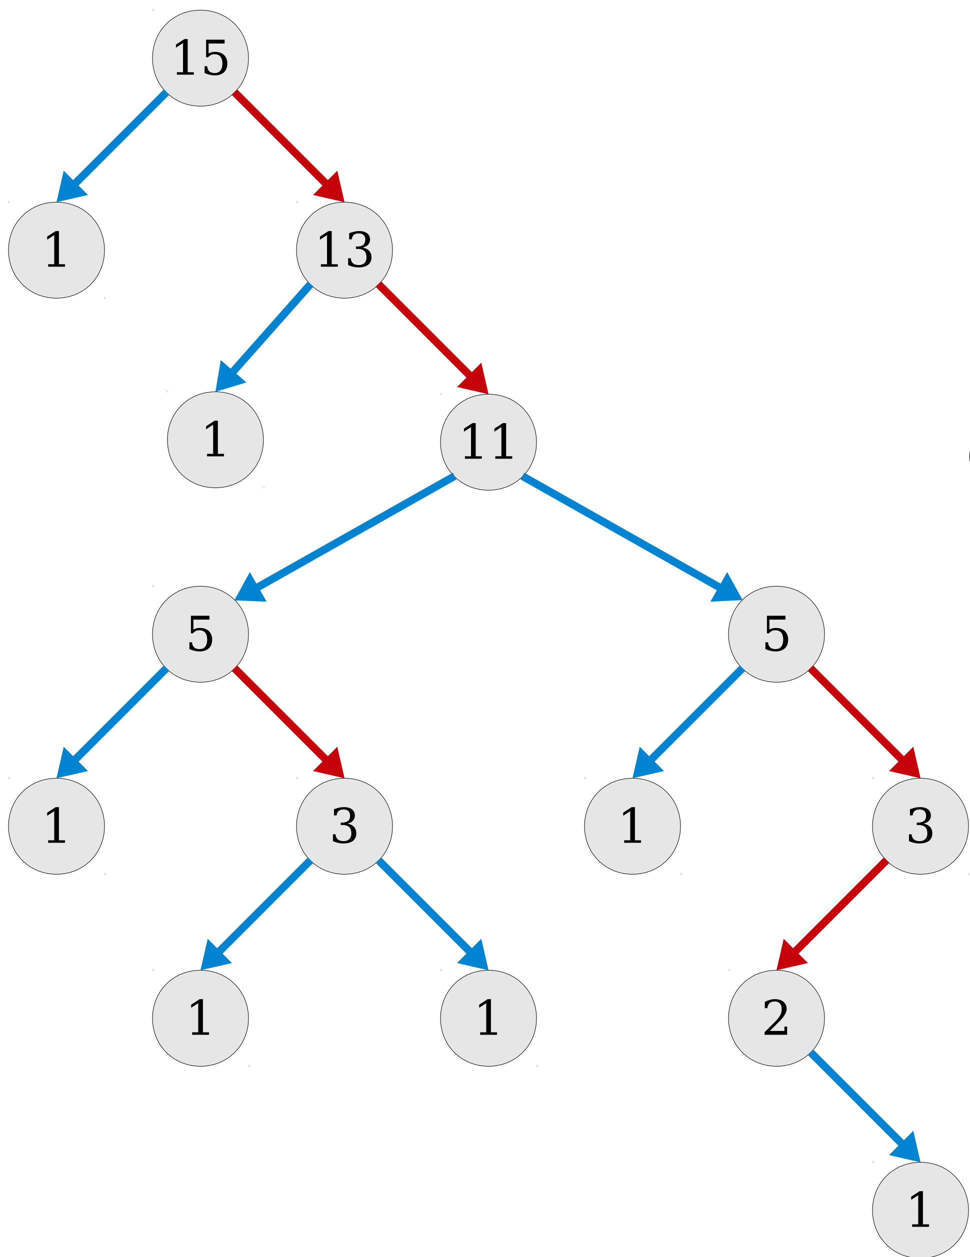

$\lg s_{new} > \lg s_{old} - 1$ (red arrow)

$\lg s_{new} \leq \lg s_{old} - 1$ (blue arrow)

Cost of looking up x_i:

O(**log (W / w_i)** + **#red-used**)

We need a potential function Φ that looks at $\lg s_i$ for each key x_i.

Reasonable guess: $\Phi = \sum_{i=1}^{n} \lg s_i$.

The Net Result

- ***Theorem:*** Using the potential function from before, the amortized cost of splaying key x_i is

$$\mathbf{1 + 3 \lg (W / w_i) = \Theta(1 + \log (W / w_i))},$$

  where $W = w_1 + w_2 + \ldots + w_n$.
- The math behind this theorem is nontrivial and not at all interesting. It's just hard math gymnastics. Check Sleator and Tarjan's paper for details!
- This theorem holds for *any* choice of weights w_i assigned to the nodes, so it's useful for proving a number of nice results about splay trees.
- There's a subtle catch, though…

An Important Detail

- ***Recall:*** When using the potential method, the sum of the amortized costs relates to the sum of the real costs as follows:

$$\sum_{i=1}^{m} a(op_i) = \sum_{i=1}^{m} t(op_i) + O(1) \cdot (\Phi_{m+1} - \Phi_1)$$

- Therefore:

$$\sum_{i=1}^{m} a(op_i) + O(1) \cdot (\Phi_1 - \Phi_{m+1}) = \sum_{i=1}^{m} t(op_i)$$

- The actual cost is bounded by the sum of the amortized costs, ***plus the drop in potential***.

An Important Detail

- Previously, when we've analyzed amortized-efficient data structures, our potential function started at 0 and ended nonnegative.
- With our current choice of

$$\Phi = \sum_{i=1}^{n} \lg s_i$$

  if we're given a tree and then start splaying, our initial potential is nonzero, and our final potential might be lower than initial.
- This means that if we perform *m* operations on an *n*-element splay tree, we need to factor $\Phi_1 - \Phi_{m+1}$ into the total cost.

Analyzing Splay Trees

- To analyze the cost of splay tree operations, we'll proceed in three steps:
  - First, assign the weights to the nodes in a way that correlates weights and access patterns.
  - Second, use the amortized cost from before to determine the cost of each splay.
  - Finally, factor in the potential drop to account for the "startup cost."
- The net result can be used to bound the cost of splaying.

Theorem (Balance Theorem): The cost of performing m operations on an n-node splay tree is $O(m \log n + n \log n)$.

Theorem (Balance Theorem): The cost of performing m operations on an n-node splay tree is $O(m \log n + n \log n)$.

Proof: The runtime of each operation is bounded by the cost of O(1) splays, so we'll bound the overall runtime by bounding the costs of the splays involved.

Theorem (Balance Theorem): The cost of performing m operations on an n-node splay tree is $O(m \log n + n \log n)$.

Proof: The runtime of each operation is bounded by the cost of O(1) splays, so we'll bound the overall runtime by bounding the costs of the splays involved.

Assign each key x_i weight $w_i = 1$.

Theorem (Balance Theorem): The cost of performing m operations on an n-node splay tree is $O(m \log n + n \log n)$.

Proof: The runtime of each operation is bounded by the cost of O(1) splays, so we'll bound the overall runtime by bounding the costs of the splays involved.

Assign each key x_i weight $w_i = 1$. The amortized cost of performing a splay at a key x_i is then

$$1 + 3 \lg (W / w_i)$$

Theorem (Balance Theorem): The cost of performing m operations on an n-node splay tree is $O(m \log n + n \log n)$.

Proof: The runtime of each operation is bounded by the cost of O(1) splays, so we'll bound the overall runtime by bounding the costs of the splays involved.

Assign each key x_i weight $w_i = 1$. The amortized cost of performing a splay at a key x_i is then

$$1 + 3 \lg (W / w_i)$$
$$= 1 + 3 \lg (n / 1)$$

Theorem (Balance Theorem): The cost of performing m operations on an n-node splay tree is $O(m \log n + n \log n)$.

Proof: The runtime of each operation is bounded by the cost of O(1) splays, so we'll bound the overall runtime by bounding the costs of the splays involved.

Assign each key x_i weight $w_i = 1$. The amortized cost of performing a splay at a key x_i is then

$$
\begin{aligned}
& 1 + 3 \lg (W / w_i) \\
&= 1 + 3 \lg (n / 1) \\
&= O(\log n).
\end{aligned}
$$

Theorem (Balance Theorem): The cost of performing m operations on an n-node splay tree is $O(m \log n + n \log n)$.

Proof: The runtime of each operation is bounded by the cost of O(1) splays, so we'll bound the overall runtime by bounding the costs of the splays involved.

Assign each key x_i weight $w_i = 1$. The amortized cost of performing a splay at a key x_i is then

$$
\begin{aligned}
& 1 + 3 \lg (W / w_i) \\
&= 1 + 3 \lg (n / 1) \\
&= O(\log n).
\end{aligned}
$$

To bound the total drop in potential, note that in the worst case any key x_i may initially be in a subtree of total weight n and end in a subtree of total weight 1.

Theorem (Balance Theorem): The cost of performing m operations on an n-node splay tree is $O(m \log n + n \log n)$.

Proof: The runtime of each operation is bounded by the cost of O(1) splays, so we'll bound the overall runtime by bounding the costs of the splays involved.

Assign each key x_i weight $w_i = 1$. The amortized cost of performing a splay at a key x_i is then

$$\begin{aligned} & 1 + 3 \lg (W / w_i) \\ = & \ 1 + 3 \lg (n / 1) \\ = & \ O(\log n). \end{aligned}$$

To bound the total drop in potential, note that in the worst case any key x_i may initially be in a subtree of total weight n and end in a subtree of total weight 1. Therefore, the maximum possible potential drop is

$$\sum_{i=1}^{n} \lg n - \sum_{i=1}^{n} \lg 1 = n \lg n.$$

Theorem (Balance Theorem): The cost of performing m operations on an n-node splay tree is $O(m \log n + n \log n)$.

Proof: The runtime of each operation is bounded by the cost of O(1) splays, so we'll bound the overall runtime by bounding the costs of the splays involved.

Assign each key x_i weight $w_i = 1$. The amortized cost of performing a splay at a key x_i is then

$$
\begin{aligned}
& 1 + 3 \lg (W / w_i) \\
&= 1 + 3 \lg (n / 1) \\
&= O(\log n).
\end{aligned}
$$

To bound the total drop in potential, note that in the worst case any key x_i may initially be in a subtree of total weight n and end in a subtree of total weight 1. Therefore, the maximum possible potential drop is

$$\sum_{i=1}^{n} \lg n - \sum_{i=1}^{n} \lg 1 = n \lg n.$$

So the total cost of performing m operations on an n-node splay tree is $O(m \log n + n \log n)$, as required. ■

A Stronger Result

- ***Recall:*** A statically optimal binary search tree has expected lookup cost $\Theta(1 + H)$, where H is the Shannon entropy of the access probability distribution.
- ***Claim:*** In a sense, splay trees achieve this statically optimal bound.

Static Optimality Theorem: Let $S = \{ x_1, ..., x_n \}$ be a set of keys stored in a splay tree. Suppose a series of lookups is performed where

- every node is accessed at least once, and
- all lookups are successful.

Then the amortized cost of each access is $O(1 + H)$, where H is the Shannon entropy of the access distribution.

Proof:

Proof: Assign each key x_i weight p_i, where p_i is the fraction of the lookups that access x_i.

Proof: Assign each key x_i weight p_i, where p_i is the fraction of the lookups that access x_i. The amortized cost of a lookup of x_i is therefore at most

$$1 + 3 \lg (W / w_i) = 1 + 3 \lg (1 / p_i) = 1 - 3 \lg p_i.$$

Proof: Assign each key x_i weight p_i, where p_i is the fraction of the lookups that access x_i. The amortized cost of a lookup of x_i is therefore at most

$$1 + 3 \lg (W / w_i) = 1 + 3 \lg (1 / p_i) = 1 - 3 \lg p_i.$$

Since each element is accessed mp_i times, the sum of the amortized lookup times is given by

$$\sum_{i=1}^{n} \left(m p_i (1 - 3 \lg p_i) \right) = m \sum_{i=1}^{n} \left(p_i - 3 p_i \lg p_i \right) = m + 3mH.$$

Proof: Assign each key x_i weight p_i, where p_i is the fraction of the lookups that access x_i. The amortized cost of a lookup of x_i is therefore at most

$$1 + 3 \lg (W / w_i) = 1 + 3 \lg (1 / p_i) = 1 - 3 \lg p_i.$$

Since each element is accessed mp_i times, the sum of the amortized lookup times is given by

$$\sum_{i=1}^{n} \left(m p_i (1 - 3 \lg p_i) \right) = m \sum_{i=1}^{n} \left(p_i - 3 p_i \lg p_i \right) = m + 3mH.$$

To bound the total drop in potential, notice that each node contributes $\lg s_i$ to the potential, where s_i is the weight of the subtree rooted at s_i.

Proof: Assign each key x_i weight p_i, where p_i is the fraction of the lookups that access x_i. The amortized cost of a lookup of x_i is therefore at most

$$1 + 3 \lg (W / w_i) = 1 + 3 \lg (1 / p_i) = 1 - 3 \lg p_i.$$

Since each element is accessed mp_i times, the sum of the amortized lookup times is given by

$$\sum_{i=1}^{n} \left(m p_i (1 - 3 \lg p_i) \right) = m \sum_{i=1}^{n} \left(p_i - 3 p_i \lg p_i \right) = m + 3 m H .$$

To bound the total drop in potential, notice that each node contributes $\lg s_i$ to the potential, where s_i is the weight of the subtree rooted at s_i. The maximum value of s_i is 1 (when all nodes are in x_i's tree) and the minimum value of s_i is p_i (when x_i is by itself), so the maximum possible potential drop from a single element is given by $-\lg p_i$.

Proof: Assign each key x_i weight p_i, where p_i is the fraction of the lookups that access x_i. The amortized cost of a lookup of x_i is therefore at most

$$1 + 3 \lg (W / w_i) = 1 + 3 \lg (1 / p_i) = 1 - 3 \lg p_i.$$

Since each element is accessed mp_i times, the sum of the amortized lookup times is given by

$$\sum_{i=1}^{n} \left(m p_i (1 - 3 \lg p_i) \right) = m \sum_{i=1}^{n} \left(p_i - 3 p_i \lg p_i \right) = m + 3mH.$$

To bound the total drop in potential, notice that each node contributes $\lg s_i$ to the potential, where s_i is the weight of the subtree rooted at s_i. The maximum value of s_i is 1 (when all nodes are in x_i's tree) and the minimum value of s_i is p_i (when x_i is by itself), so the maximum possible potential drop from a single element is given by $-\lg p_i$. Therefore, the maximum potential drop is

$$\sum_{i=1}^{n} -\lg p_i \le \sum_{i=1}^{n} -m p_i \lg p_i = m \sum_{i=1}^{n} -p_i \lg p_i = mH$$

Proof: Assign each key x_i weight p_i, where p_i is the fraction of the lookups that access x_i. The amortized cost of a lookup of x_i is therefore at most

$$1 + 3 \lg (W / w_i) = 1 + 3 \lg (1 / p_i) = 1 - 3 \lg p_i.$$

Since each element is accessed mp_i times, the sum of the amortized lookup times is given by

$$\sum_{i=1}^{n} \left(m p_i (1 - 3 \lg p_i)\right) = m \sum_{i=1}^{n} \left(p_i - 3 p_i \lg p_i\right) = m + 3mH.$$

To bound the total drop in potential, notice that each node contributes $\lg s_i$ to the potential, where s_i is the weight of the subtree rooted at s_i. The maximum value of s_i is 1 (when all nodes are in x_i's tree) and the minimum value of s_i is p_i (when x_i is by itself), so the maximum possible potential drop from a single element is given by $-\lg p_i$. Therefore, the maximum potential drop is

$$\sum_{i=1}^{n} -\lg p_i \leq \sum_{i=1}^{n} -m p_i \lg p_i = m \sum_{i=1}^{n} -p_i \lg p_i = mH$$

So the cost of the m lookups is $O(m + mH)$, and since there are m lookups, the amortized cost of each is $O(1 + H)$. ■

Beating Static Optimality

- On many classes of access sequences, splay trees can outperform statically optimal BSTs.
- The ***sequential access theorem*** says that

  If you look up all n elements in a splay tree in ascending order, the amortized cost of each lookup is $O(1)$.

- The ***working-set theorem*** says that

  If you perform $\Omega(n \log n)$ successful lookups, the amortized cost of each successful lookup is $O(1 + \log t)$, where t is the number of searches since we last looked up the element searched for.

- In the upcoming programming assignment, you'll compare the performance of splay trees to (nearly) optimal BSTs. See if you notice anything interesting in these cases!

An Open Problem: ***Dynamic Optimality***

Many Flavors of BSTs

- Over the course of this quarter, we've seen a bunch of different types of BSTs:
  - Weight-balanced trees (PS2, PS4)
  - Red/black trees
  - AVL trees
  - Splay trees
- Each tree structure makes a different set of tradeoffs (per-operation efficient vs. amortized efficient, fast lookups vs. fast insertions, static vs. dynamic, etc.)
- ***Question:*** Is there a single BST data structure that's the "best" possible choice across all BSTs?

The Offline Perfect BST

- Imagine that you're told, in advance, that you'll be maintaining a binary search tree that will have a particular series of operations X performed on it.
- You're allowed unlimited time to plan out the exact sequence of rotations and updates you're going to make on your tree.
- The cost of your solution is the number of primitive tree operations performed on your BST (for example, following pointers, performing rotations, etc.)
- We'll denote by ***OPT(X)*** the minimum possible cost associated with performing your series of operations on your BST.

Dynamic Optimality

- Suppose you have a binary search tree data structure T maintained according to some algorithm (e.g. red/black rules, splaying, etc.)
- We say that T is ***dynamically optimal*** if the cost of performing *any* series of operations X on T is $O(OPT(X))$.
- For example, red/black trees are *not* dynamically optimal, since the cost of performing a low-entropy series of lookups on an optimal BST is $O(m + mH)$, whereas it could be $\Theta(m \log n)$ in a red/black tree.

The Conjecture

- ***Conjectured:*** Splay trees are dynamically optimal.
    - In other words, there is no known series of operations we can perform on a splay tree that does any more than a constant factor worse than the cost of performing those operations on *any* dynamic search tree!
- We're probably still a ways off on being able to prove or disprove this statement. It's an active area of research!

More to Explore

- In 2004, Demaine et al. invented the ***tango tree***, which is at most an O(log log n) factor away from dynamic optimality.
- In 2006, Wang et al. developed the ***multisplay tree***, which is also at most a factor of O(log log n) from dynamic optimality.
- In 2007, Pettie proved that if a splay tree is used to implement a deque, then the amortized cost of each operation is **$O(\alpha^*(n))$**, where $\alpha^*(n)$ is the number of times the Ackermann inverse function needs to be applied to n to drop it to a constant.
- In 2009, Demaine et al. published "***The Geometry of Binary Search Trees***," providing a new framework for analyzing dynamic optimality.
- In 2012, Bose et al. developed the ***crazy good chocolate pop tart***, a type of stack implemented with a binary search tree, and used it to build a worst-case efficient version of the splay tree.

Next Time

- ***Randomized Data Structures***
  - How do we trade worst-case guarantees for probabilistic guarantees?
- ***Count[-Min] Sketches***
  - Counting in sublinear space.
- ***Concentration Inequalities***
  - How do we show that we're near the expected value most of the time?