van Emde Boas Trees

Outline for Today

- **Data Structures on Integers**
  - How can we speed up operations that work on integer data?

- **Tiered Bitvectors**
  - A simple data structure for ordered dictionaries.

- **van Emde Boas Trees**
  - An extremely fast data structure for ordered dictionaries.

Integer Data Structures

Working with Integers

- Integers are interesting objects to work with:
  - They can be treated as strings of bits, so we can use techniques from string processing.
  - They fit into machine words, so we can process the bits in parallel with individual word operations.
- Today, we'll explore **van Emde Boas trees**, which rely on this second property.
- Wednesday, we'll see **y-fast tries**, which will pull together just about everything from the quarter.

Our Machine Model

- We will assume that we are working with a ***transdichotomous machine model***.

- Memory is split apart into integer words composed of w bits each.

- The CPU can perform basic arithmetic operations (addition, subtraction, multiplication, division, shifts, AND, OR, etc.) on machine words in time O(1) each.

- When working on a problem where each instance has size n, we assume $w = \Omega(\log n)$.

Ordered Dictionaries

Ordered Dictionaries

- An **ordered dictionary** is a data structure that maintains a set S of elements drawn from an ordered universe $\mathscr{U}$ and supports these operations:

  - **insert**(x), which adds x to S.

  - **is-empty**(), which returns whether $S = \emptyset$.

  - **lookup**(x), which returns whether $x \in S$.

  - **delete**(x), which removes x from S.

  - **max**() / **min**(), which returns the maximum or minimum element of S.

  - **successor**(x), which returns the smallest element of S greater than x, and

  - **predecessor**(x), which returns the largest element of S smaller than x.

Ordered Dictionaries

- Balanced BSTs support all ordered dictionary operations in time O(log n) each.

- Hash tables support insertion, lookups, and deletion in expected time O(1), but require time O(n) for min, max, successor, and predecessor.

Ordered Integer Dictionaries

- Suppose that our universe consists of natural numbers upper-bounded by some number U.

  - Specifically, $\mathscr{U} = [U] = \{0, 1, 2, ..., U - 1\}$.

- **Question:** Can we design a data structure that supports the ordered dictionary operations on $\mathscr{U}$ faster than a balanced BST?

- The answer is *yes*, and we'll see van Emde Boas trees and *y*-fast tries as two possible solutions.

A Preliminary Approach: **Bitvectors**

Bitvectors

- A **bitvector** is an array of bits of length U.

- Represents a set of elements with O(1) insertions, deletions, and lookups:

  - To insert x, set the bit for x to 1.

  - To delete x, set the bit for x to 0.

  - To lookup x, check whether the bit for x is 1.

- Space usage is $\Theta(U)$.

110111001011101111000100110101011110011011110111

Bitvectors

- A **bitvector** is an array of bits of length U.

- Represents a set of elements with O(1) insertions, deletions, and lookups:

  - To insert x, set the bit for x to 1.

  - To delete x, set the bit for x to 0.

  - To lookup x, check whether the bit for x is 1.

- Space usage is $\Theta(U)$.

110111100101011101111000100110101010111100110111101111

Bitvectors

- A **bitvector** is an array of bits of length U.

- Represents a set of elements with O(1) insertions, deletions, and lookups:

  - To insert x, set the bit for x to 1.

  - To delete x, set the bit for x to 0.

  - To lookup x, check whether the bit for x is 1.

- Space usage is $\Theta(U)$.

110110001011101111000100110101011110011011101111

Bitvectors

- A **bitvector** is an array of bits of length U.

- Represents a set of elements with O(1) insertions, deletions, and lookups:

  - To insert x, set the bit for x to 1.

  - To delete x, set the bit for x to 0.

  - To lookup x, check whether the bit for x is 1.

- Space usage is $\Theta(U)$.

110110001011101111000100110101011110011011110111 1

Bitvectors

- A **bitvector** is an array of bits of length U.

- Represents a set of elements with $O(1)$ insertions, deletions, and lookups:

  - To insert x, set the bit for x to 1.

  - To delete x, set the bit for x to 0.

  - To lookup x, check whether the bit for x is 1.

- Space usage is $\Theta(U)$.

11011000101110111100010011010101111001101111011111

Bitvectors

- A **bitvector** is an array of bits of length U.

- Represents a set of elements with O(1) insertions, deletions, and lookups:

  - To insert x, set the bit for x to 1.

  - To delete x, set the bit for x to 0.

  - To lookup x, check whether the bit for x is 1.

- Space usage is $\Theta(U)$.

11011000101110111101010011010101111001101111101111

Bitvectors

- A **bitvector** is an array of bits of length U.

- Represents a set of elements with O(1) insertions, deletions, and lookups:

  - To insert x, set the bit for x to 1.

  - To delete x, set the bit for x to 0.

  - To lookup x, check whether the bit for x is 1.

- Space usage is $\Theta(U)$.

110110001011101111010100110101011110011011110111

Bitvectors

- The min, max, predecessor, and successor operations on bitvectors can be extremely slow.

- Runtime will be $\Theta(U)$ in the worst case.

```
00000000000000000000000001000000000000000000000000000
```

Tiered Bitvectors

- Adapting an approach similar to our hybrid RMQs, we can put a summary structure on top of our bitvector.

- Break the universe U into $\Theta(U / B)$ blocks of size B.

- Create an auxiliary bitvector of size $\Theta(U / B)$ that stores which blocks are nonempty.

`00100010000000000011000000000000000010011110111000000000000000000`

Tiered Bitvectors

- Adapting an approach similar to our hybrid RMQs, we can put a summary structure on top of our bitvector.

- Break the universe U into $\Theta(U\ /\ B)$ blocks of size B.

- Create an auxiliary bitvector of size $\Theta(U\ /\ B)$ that stores which blocks are nonempty.

```
00100010 00000000 00011000 00000000 00000100 11110111 00000000 00000000
```

Tiered Bitvectors

- Adapting an approach similar to our hybrid RMQs, we can put a summary structure on top of our bitvector.

- Break the universe U into $\Theta(U / B)$ blocks of size B.

- Create an auxiliary bitvector of size $\Theta(U / B)$ that stores which blocks are nonempty.

00100010	00000000	00011000	00000000	00000100	11110111	00000000	00000000

Tiered Bitvectors

- Adapting an approach similar to our hybrid RMQs, we can put a summary structure on top of our bitvector.

- Break the universe U into $\Theta(U / B)$ blocks of size B.

- Create an auxiliary bitvector of size $\Theta(U / B)$ that stores which blocks are nonempty.

1	0	1	0	1	1	0	0
00100010	00000000	00011000	00000000	00000100	11110111	00000000	00000000

Tiered Bitvectors

- Using the same techniques we used for RMQ, we can speed up ordered dictionary operations so that they run in time $O(U / B + B)$.

- As before, this is minimized when $B = \Theta(U^{1/2})$.

- Ordered dictionary runtimes are now all $O(U^{1/2})$.

1	0	1	0	1	1	0	0
00100010	00000000	00011000	00000000	00000100	11110111	00000000	00000000

Tiered Bitvectors

- This approach does introduce some slowdown to the *delete* operation.

- Whenever we do a *delete*, we have to check whether the block is now empty and, if so, we need to clear the bit in the summary bitvector.

- New cost: $O(U^{1/2})$.

1	0	1	0	1	1	0	0
00100010	00000000	00011000	00000000	00000100	11110111	00000000	00000000

Tiered Bitvectors

- This approach does introduce some slowdown to the ***delete*** operation.

- Whenever we do a ***delete***, we have to check whether the block is now empty and, if so, we need to clear the bit in the summary bitvector.

- New cost: $O(U^{1/2})$.

1	0	1	0	1	1	0	0
00100010	00000000	0001**1**000	00000000	00000100	11110111	00000000	00000000

Tiered Bitvectors

- This approach does introduce some slowdown to the *delete* operation.

- Whenever we do a *delete*, we have to check whether the block is now empty and, if so, we need to clear the bit in the summary bitvector.

- New cost: $O(U^{1/2})$.

1	0	1	0	1	1	0	0
00100010	00000000	00010000	00000000	00000100	11110111	00000000	00000000

Tiered Bitvectors

- This approach does introduce some slowdown to the *delete* operation.

- Whenever we do a *delete*, we have to check whether the block is now empty and, if so, we need to clear the bit in the summary bitvector.

- New cost: $O(U^{1/2})$.

1	0	1	0	1	1	0	0
00100010	00000000	00010000	00000000	00000100	11110111	00000000	00000000

Tiered Bitvectors

- This approach does introduce some slowdown to the ***delete*** operation.

- Whenever we do a ***delete***, we have to check whether the block is now empty and, if so, we need to clear the bit in the summary bitvector.

- New cost: $O(U^{1/2})$.

1	0	1	0	1	1	0	0
00100010	00000000	000**1**0000	00000000	00000100	11110111	00000000	00000000

Tiered Bitvectors

- This approach does introduce some slowdown to the *delete* operation.

- Whenever we do a *delete*, we have to check whether the block is now empty and, if so, we need to clear the bit in the summary bitvector.

- New cost: $O(U^{1/2})$.

1	0	1	0	1	1	0	0
00100010	00000000	00000000	00000000	00000100	11110111	00000000	00000000

Tiered Bitvectors

- This approach does introduce some slowdown to the **_delete_** operation.

- Whenever we do a **_delete_**, we have to check whether the block is now empty and, if so, we need to clear the bit in the summary bitvector.

- New cost: O($U^{1/2}$).

1	0	0	0	1	1	0	0
00100010	00000000	0000**0**0000	00000000	00000100	11110111	00000000	00000000

Tiered Bitvectors

- This approach does introduce some slowdown to the *delete* operation.

- Whenever we do a *delete*, we have to check whether the block is now empty and, if so, we need to clear the bit in the summary bitvector.

- New cost: $O(U^{1/2})$.

1	0	0	0	1	1	0	0
00100010	00000000	00000000	00000000	00000100	11110111	00000000	00000000

Tiered Bitvectors

- We can view our tiered bitvector structure in a different light that will help lead to future improvements.

- Instead of thinking of this as two bitvectors (a main and a summary), think of it as $\Theta(U^{1/2})$ smaller main bitvectors and a summary bitvector.

1	0	0	0	1	1	0	0
00100010	00000000	00000000	00000000	00000100	11110111	00000000	00000000

Tiered Bitvectors

- To perform **_lookup_**(x) in this structure, check the $\lfloor x / U^{1/2} \rfloor$th bitvector to see if $x \bmod U^{1/2}$ is present.

- In other words, our top-level **_lookup_**(x) call turns into a recursive **_lookup_**($\lfloor x / U^{1/2} \rfloor$) call in a smaller bitvector.

1	0	0	0	1	1	0	0
00100010	00000000	00000000	00000000	00000100	11110111	00000000	00000000

Tiered Bitvectors

- To perform **lookup**(x) in this structure, check the $\lfloor x / U^{1/2} \rfloor$th bitvector to see if $x \bmod U^{1/2}$ is present.

- In other words, our top-level **lookup**(x) call turns into a recursive **lookup**($\lfloor x / U^{1/2} \rfloor$) call in a smaller bitvector.

1	0	0	0	1	1	0	0
00100010	00000000	00000000	00000000	00000100	11110111	00000000	00000000

$$x = 42$$

Tiered Bitvectors

- To perform **lookup**(x) in this structure, check the $\lfloor x / U^{1/2} \rfloor$th bitvector to see if $x \bmod U^{1/2}$ is present.

- In other words, our top-level **lookup**(x) call turns into a recursive **lookup**($\lfloor x / U^{1/2} \rfloor$) call in a smaller bitvector.

1	0	0	0	1	1	0	0
00100010	00000000	00000000	00000000	00000100	11110111	00000000	00000000

$\lfloor 42 / 8 \rfloor = 5$ $42 \bmod 8 = 2$

Tiered Bitvectors

- To perform **lookup**(x) in this structure, check the $\lfloor x / U^{1/2} \rfloor$th bitvector to see if $x \bmod U^{1/2}$ is present.

- In other words, our top-level **lookup**(x) call turns into a recursive **lookup**($\lfloor x / U^{1/2} \rfloor$) call in a smaller bitvector.

1	0	0	0	1	1	0	0
00100010	00000000	00000000	00000000	00000100	11110111	00000000	00000000

$$\lfloor 42 / 8 \rfloor = 5 \qquad 42 \bmod 8 = 2$$

Tiered Bitvectors

- To perform ***insert***(x) in this structure, insert $x \bmod U^{1/2}$ into the $\lfloor x \,/\, U^{1/2} \rfloor$th bitvector, then insert $\lfloor x \,/\, U^{1/2} \rfloor$ into the summary bitvector.

- Turns one ***insert*** call into two recursive ***insert*** calls.

1	0	0	0	1	1	0	0
00100010	00000000	00000000	00000000	00000100	11110111	00000000	00000000

Tiered Bitvectors

- To perform **_insert_**(x) in this structure, insert $x \bmod U^{1/2}$ into the $\lfloor x / U^{1/2} \rfloor$th bitvector, then insert $\lfloor x / U^{1/2} \rfloor$ into the summary bitvector.

- Turns one **_insert_** call into two recursive **_insert_** calls.

1	0	0	0	1	1	0	0
00100010	00000000	00000000	00000000	00000100	11110111	00000000	00000000

$x = 25$

Tiered Bitvectors

- To perform $\textbf{\textit{insert}}(x)$ in this structure, insert $x \bmod U^{1/2}$ into the $\lfloor x / U^{1/2} \rfloor$th bitvector, then insert $\lfloor x / U^{1/2} \rfloor$ into the summary bitvector.

- Turns one $\textbf{\textit{insert}}$ call into two recursive $\textbf{\textit{insert}}$ calls.

1	0	0	0	1	1	0	0
00100010	00000000	00000000	00000000	00000100	11110111	00000000	00000000

$$\lfloor 25 / 8 \rfloor = 3 \qquad 25 \bmod 8 = 1$$

Tiered Bitvectors

- To perform **_insert_**(x) in this structure, insert $x \bmod U^{1/2}$ into the $\lfloor x / U^{1/2} \rfloor$th bitvector, then insert $\lfloor x / U^{1/2} \rfloor$ into the summary bitvector.

- Turns one **_insert_** call into two recursive **_insert_** calls.

1	0	0	0	1	1	0	0
00100010	00000000	00000000	00000000	00000100	11110111	00000000	00000000

$$\lfloor 25 / 8 \rfloor = 3 \qquad 25 \bmod 8 = 1$$

Tiered Bitvectors

- To perform ***insert***(x) in this structure, insert $x \bmod U^{1/2}$ into the $\lfloor x / U^{1/2} \rfloor$th bitvector, then insert $\lfloor x / U^{1/2} \rfloor$ into the summary bitvector.

- Turns one ***insert*** call into two recursive ***insert*** calls.

1	0	0	1	1	1	0	0
00100010	00000000	00000000	01000000	00000100	11110111	00000000	00000000

$$\lfloor 25 / 8 \rfloor = 3 \qquad 25 \bmod 8 = 1$$

Tiered Bitvectors

- To perform *max*(), call *max* on the summary structure.

- If it returns value v, return *max* of the vth bitvector.

- Turns one *max* call into two recursive *max*s.

1	0	0	0	1	1	0	0
00100010	00000000	00000000	00000000	00000100	11110111	00000000	00000000

Tiered Bitvectors

- To perform $max()$, call max on the summary structure.

- If it returns value v, return max of the vth bitvector.

- Turns one max call into two recursive maxs.

1	0	0	0	1	1	0	0
00100010	00000000	00000000	00000000	00000100	11110111	00000000	00000000

Tiered Bitvectors

- To perform *max*(), call *max* on the summary structure.

- If it returns value v, return *max* of the vth bitvector.

- Turns one *max* call into two recursive *max*s.

1	0	0	0	1	1	0	0
00100010	00000000	00000000	00000000	00000100	11110111	00000000	00000000

Tiered Bitvectors

- To perform **successor**(*x*), do the following:

  - Find **max** in the $\lfloor x / U^{1/2} \rfloor$th bitvector.

  - If it exists and is greater than *x*, find **successor**(*x* mod $U^{1/2}$) in that bitvector.

  - Otherwise, find **successor**($\lfloor x / U^{1/2} \rfloor$) in the summary structure; let it be *j* if it exists.

  - Return **min** of the *j*th bitvector of it exists or ∞ otherwise.

- Turns **successor** into a **max**, a **min**, and a **successor**.

1	0	0	0	1	1	0	0
00100010	00000000	00000000	00000000	00000100	11110111	00000000	00000000

Tiered Bitvectors

- To perform an *is-empty* query, return the result of that query on the summary structure.

- Turns one *is-empty* query into a single smaller *is-empty* query.

1	0	0	0	1	1	0	0
00100010	00000000	00000000	00000000	00000100	11110111	00000000	00000000

Tiered Bitvectors

- To perform **delete**(x) in this structure, delete $x \bmod U^{1/2}$ from the $\lfloor x / U^{1/2} \rfloor$th bitvector.

- Then, check **is-empty** on that bitvector, and if so, **delete**($\lfloor x / U^{1/2} \rfloor$) from the summary bitvector.

- Turns one **delete** call into up to two recursive **delete**s and one **is-empty**.

1	0	0	0	1	1	0	0
00100010	00000000	00000000	00000000	00000100	11110111	00000000	00000000

Tiered Bitvectors

- To perform **_delete_**(x) in this structure, delete x mod $U^{1/2}$ from the $\lfloor x / U^{1/2} \rfloor$th bitvector.

- Then, check **_is-empty_** on that bitvector, and if so, **_delete_**($\lfloor x / U^{1/2} \rfloor$) from the summary bitvector.

- Turns one **_delete_** call into up to two recursive **_delete_**s and one **_is-empty_**.

1	0	0	0	1	1	0	0
00100010	00000000	00000000	00000000	00000**1**00	11110111	00000000	00000000

Tiered Bitvectors

- To perform ***delete***(x) in this structure, delete $x \bmod U^{1/2}$ from the $\lfloor x / U^{1/2} \rfloor$th bitvector.

- Then, check ***is-empty*** on that bitvector, and if so, ***delete***$(\lfloor x / U^{1/2} \rfloor)$ from the summary bitvector.

- Turns one ***delete*** call into up to two recursive ***delete***s and one ***is-empty***.

1	0	0	0	0	1	0	0
00100010	00000000	00000000	00000000	00000000	11110111	00000000	00000000

The Story So Far

- Each operation turns into recursive operations on a smaller bitvector:
  - *insert*: 2x *insert*
  - *lookup*: 1x *lookup*
  - *is-empty*: 1x *is-empty*
  - *min*: 2x *min*
  - *successor*: 1x *successor*, 1x *max*, 1x *min*
  - *delete*: 2x *delete*, 1x *is-empty*

A Recursive Approach

- Adding one tier to the bitvector sped things up appreciably.

- **Idea:** What if we apply this same approach to each of the smaller bitvectors?

- Builds a recursive data structure:

  - If $U \leq 2$, just use a normal bitvector.

  - Otherwise, split the input apart into $\Theta(U^{1/2})$ blocks of size $\Theta(U^{1/2})$ and add a summary data structure on top.

  - Answer queries using the recursive structure from before.

Our Data Structure

- Let $\mathscr{U} = [256]$.

- The top-level structure looks like this:
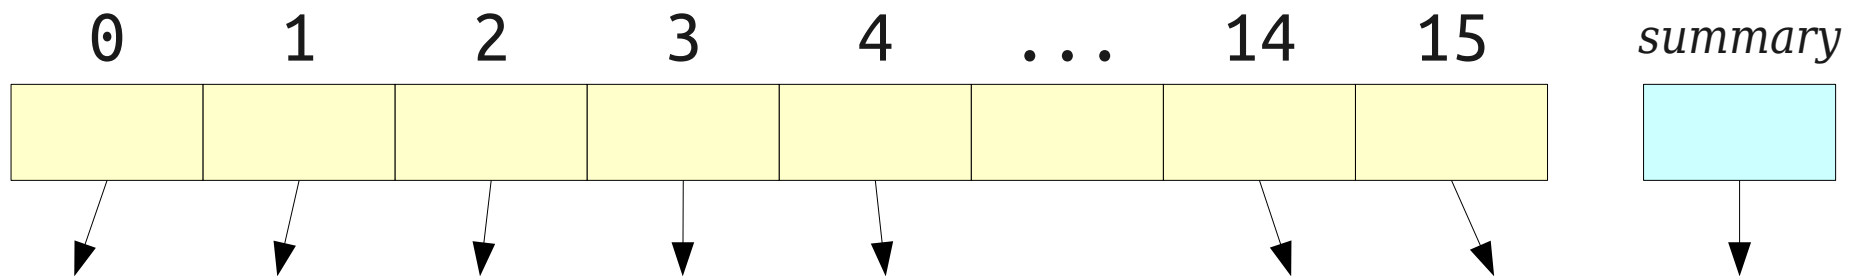
- Each structure one level below (and the summary) looks like this:
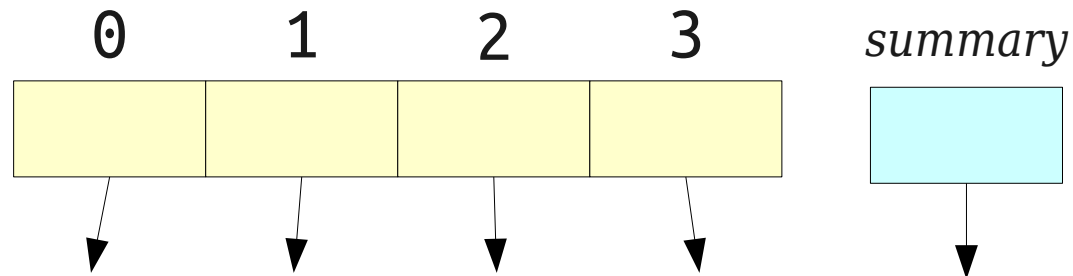
So... how efficient is it?

Analyzing the Operations

- Let's analyze the ***is-empty*** and ***lookup*** operations in this structure.

- Each makes a recursive call to a problem of size $\Theta(U^{1/2})$ and does O(1) work.

- Recurrence relation:

$$T(2) = \Theta(1)$$
$$T(U) \leq T(U^{1/2}) + \Theta(1)$$

- How do we solve this recurrence?

A Useful Substitution

- The Master Theorem is great for working with recurrences of the form

$$T(n) \leq aT(n / b) + \mathrm{O}(n^d)$$

- This recurrence doesn't have this form because the "shrinking" step is a square root rather than a division.

- To address this, we'll transform the recurrence so that it fits into the above form.

- If we write $U = 2^k$, then $U^{1/2} = 2^{k/2}$.

- Turn the recurrence from a recurrence in U to a recurrence in $k = \log U$.

The Substitution

- Define $S(k) = T(2^k)$.

- Since

$$T(2) \leq \Theta(1)$$
$$T(U) \leq T(U^{1/2}) + \Theta(1)$$

- We have

$$S(1) \leq \Theta(1)$$
$$S(k) \leq S(k \: / \: 2) + \Theta(1)$$

- This means that $S(k) = O(\log k)$.

- So $T(U) = T(2^{\lg U}) = S(\lg U) = \textbf{O(log log } U\textbf{)}$.

Analyzing the Operations

- The **_insert_** and **_min_** operations each make two recursive calls on subproblems of size $\Theta(U^{1/2})$ and do $\Theta(1)$ work.

- Gives this recurrence:

$$T(2) \leq \Theta(1)$$
$$T(U) \leq 2T(U^{1/2}) + \Theta(1)$$

- Substituting $S(k) = T(2^k)$ yields

$$S(1) \leq \Theta(1)$$
$$S(k) \leq 2S(k / 2) + \Theta(1)$$

- So $S(k) = O(k)$.

- Therefore, $T(U) = S(2^{\lg U}) = $ **O(log U)**.

Analyzing the Operations

- Each **delete** call makes two recursive **delete** calls and one call to **is-empty**.

- As we saw, **is-empty** takes time $O(\log \log U)$

- Recurrence relation is

$$T(2) \leq \Theta(1)$$
$$T(U) \leq 2T(U^{1/2}) + O(\log \log U)$$

- Letting $S(k) = T(2^k)$ gives

$$S(1) \leq \Theta(1)$$
$$S(k) \leq 2S(k / 2) + O(\log k)$$

- Via the Master Theorem, $S(k) = O(k)$.

- Thus $T(U) = \mathbf{O(\log\ U)}$.

Analyzing the Operations

- Each **_successor_** call makes one recursive **_successor_** call and one call to **_max_** and **_min_**.

- As we saw, **_max_** and **_min_** takes time O(log U)

- Recurrence relation is

$$T(2) \leq \Theta(1)$$
$$T(U) \leq T(U^{1/2}) + O(\log U)$$

- Letting $S(k) = T(2^k)$ gives

$$S(1) \leq \Theta(1)$$
$$S(k) \leq T(k/2) + O(k)$$

- Via the Master Theorem, $S(k) = O(k)$.

- Thus $T(U) = $ **O(log U)**.

Where We Stand

- Right now, we have a data structure where lookups are exponentially faster than a balanced BST if $n = \Omega(\log U)$.

- Other operations have runtime proportional to $\log U$, which is (usually) greater than $\log n$.

- ***Can we speed things up?***

Time-Out for Announcements!

Midterm Logistics

- The midterm is this Wednesday from 7PM – 10PM.
- Rooms assigned by last name:
  - A – S: Go to **Meyer Forum** (Meyer 124)
  - T – Z: Go to **Meyer 147**.
- You can bring a double-sided sheet of 8.5" × 11" paper with any notes you would like.
- Any topics up through and including today's lecture may be covered.
- Review session tonight from 7:30PM – 9:30PM in Gates 104.
- We'll hold an alternate exam from **4PM – 7PM** in **Gates 159** on the exam day; please email us ASAP if you would like to take the exam at this time.

Final Project Topics

- Approximate Distance Oracles

- Binary Decision Diagrams

- Burrows-Wheeler Transforms

- Cardinality Estimators

- Deterministic Skip Lists

- Extensible Hashing

- Hopscotch Hashing

- Link/Cut Trees

- Lock-Free Queues

- Nearest-Neighbor Searching

- R-Trees

- Robin Hood Hashing

- Ropes

- Scapegoat Trees

- Segment Trees

- Soft Heaps

Final Project Presentations

- Final project presentations will run during Week 10.

- We'll send out a signup form at 5:45PM tonight.

- Please have one person from your group choose a time slot and list the names of your group members.

- Time slot choices are final – please make sure you can make the time you choose!

Your Questions!

"When designing data structures, how do you know what is "good" and what is not? That is, sometimes we are happy with linear, sometimes with logarithmic, etc. How do we know how good we should aim for?"

It's really on a case-by-case basis. In some cases, "good" might mean "anything better than the naïve approach." In other areas where there's more progress, it might be "better than the current best solution." When lower bounds exist, it might be "matching the lower bound." In practical settings, it can mean "fast enough to work on large inputs."

"What is your favorite proof? What was your favorite data structure to code up?"

That's a tough one!

The proof of Cantor's theorem is simple, straightforward, and totally counterintuitive. It's one of my favorites since it blows everyones' minds the first time they see it.

I think the most fun I had with a data structure was with the binomial heap, since it was so much fun watching the theory I'd read in CLRS actually work out. Plus, it was fun getting to implement binary arithmetic!

"If you were taking a midterm like this one, what would you put on your one-page cheat sheet?"

I don't think I can really answer honestly since I know what's on the exam. ☺

I'd probably write out a summary of all the main data structures and the key tricks, then review it a day later and write down all the topics I couldn't fully remember. The act of writing things out really helps some people (like me!) learn things.

"What's the best way to receive one-on-one help with problem sets (perhaps pertaining to a specific aspect of a solution attempt)? Office hours sometimes don't work if there's a crowd."

You can always email us with questions if you'd like. If you'd like to meet one-on-one with us, send us an email and we can try to work something out!

Back to CS166!

Identifying Inefficiencies

- A few operations seem like easy candidates for speedups:

  - *is-empty* certainly seems like it shouldn't take time O(log log U).

  - *max* and *min* can probably don't actually need time O(log U).

- We'll show how to speed up these three operations.

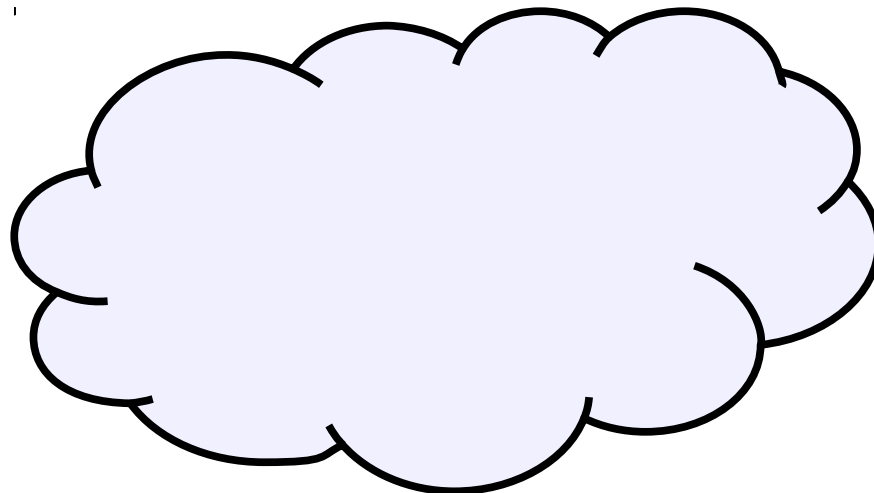- By doing so, we'll significantly improve the runtimes of the other operations.

Improving Min and Max

- Suppose you have a priority queue where finding the min takes time $\omega(1)$.

- How could you modify it so that finding the min can be done in time $O(1)$?

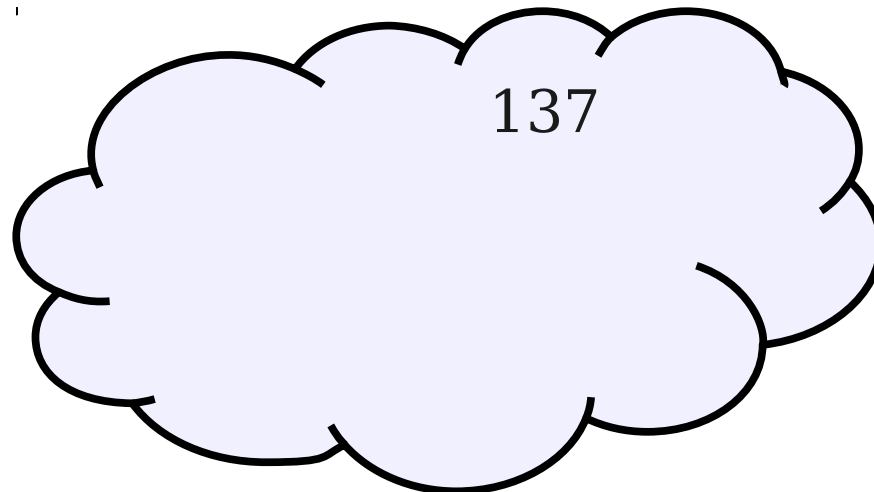- **Answer:** Store the minimum outside of the priority queue.

min

Improving Min and Max

- Suppose you have a priority queue where finding the min takes time $\omega(1)$.

- How could you modify it so that finding the min can be done in time $O(1)$?

- **Answer:** Store the minimum outside of the priority queue.

137

min

Improving Min and Max

- Suppose you have a priority queue where finding the min takes time ω(1).

- How could you modify it so that finding the min can be done in time O(1)?

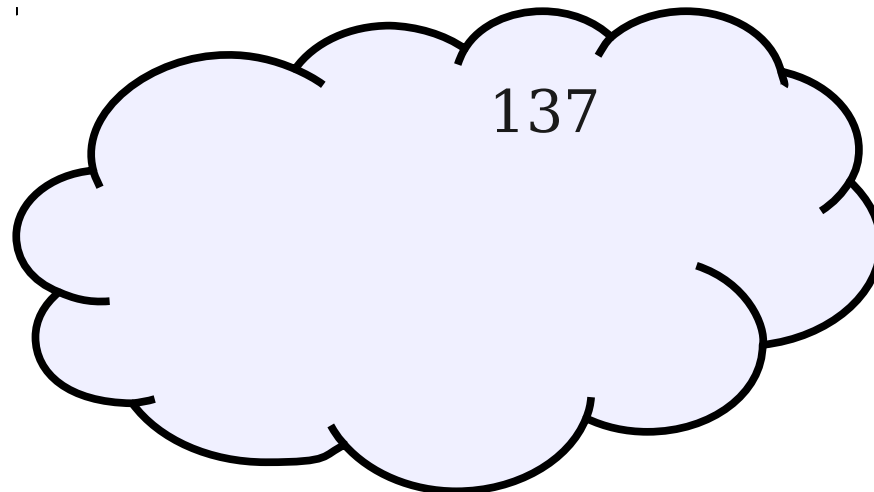- **Answer:** Store the minimum outside of the priority queue.

137

min

Improving Min and Max

- Suppose you have a priority queue where finding the min takes time $\omega(1)$.

- How could you modify it so that finding the min can be done in time $O(1)$?

- **Answer:** Store the minimum outside of the priority queue.
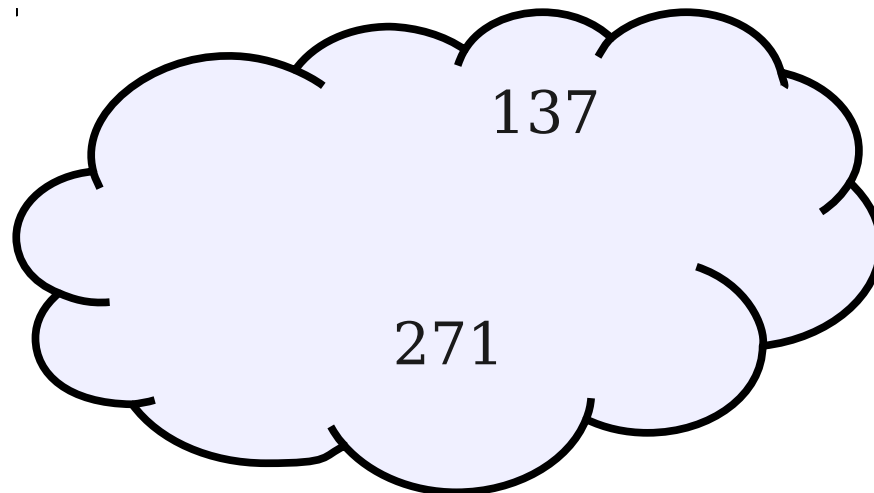
42   137

min

Improving Min and Max

- Suppose you have a priority queue where finding the min takes time $\omega(1)$.

- How could you modify it so that finding the min can be done in time $O(1)$?

- **Answer:** Store the minimum outside of the priority queue.

42

min

137

Improving Min and Max

- Suppose you have a priority queue where finding the min takes time ω(1).

- How could you modify it so that finding the min can be done in time O(1)?

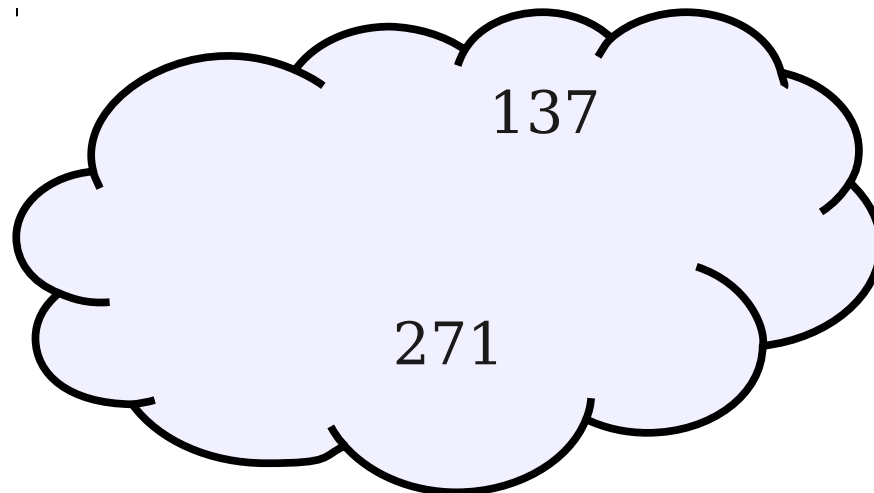- **Answer:** Store the minimum outside of the priority queue.

271

42

min

137

Improving Min and Max

- Suppose you have a priority queue where finding the min takes time ω(1).

- How could you modify it so that finding the min can be done in time O(1)?

- **Answer:** Store the minimum outside of the priority queue.

42

min

137

271

Improving Min and Max

- Suppose you have a priority queue where finding the min takes time $\omega(1)$.

- How could you modify it so that finding the min can be done in time $O(1)$?

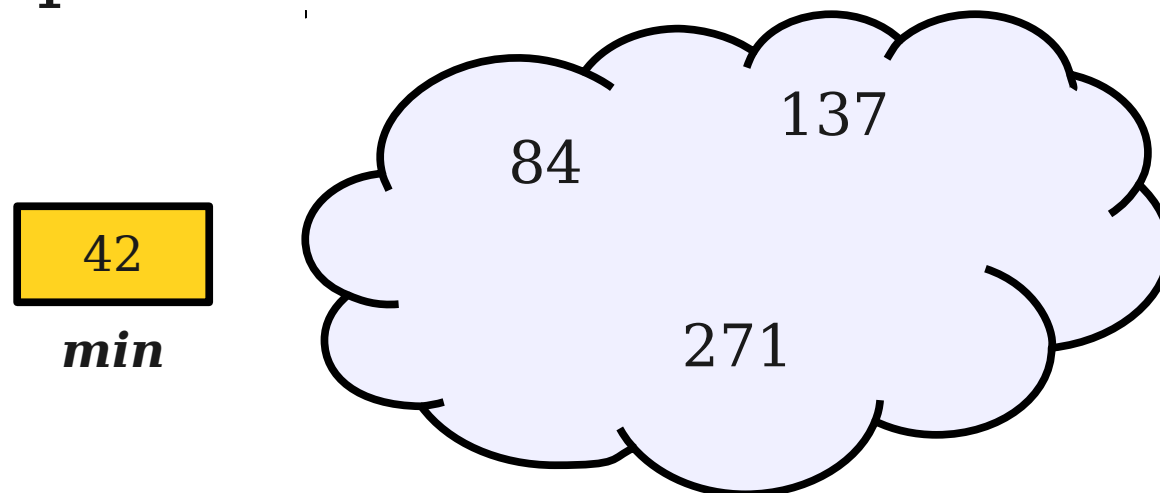- **Answer:** Store the minimum outside of the priority queue.
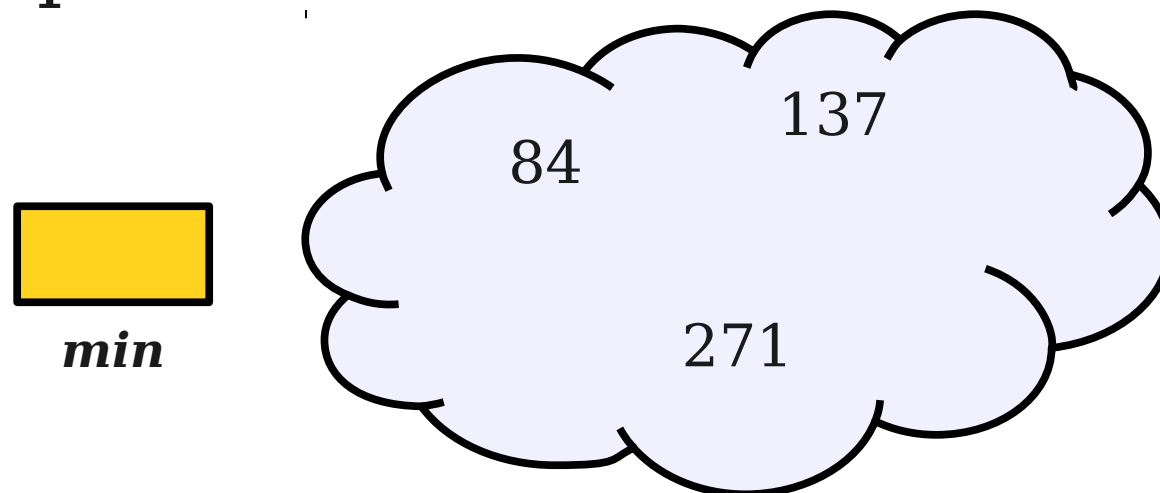
84

42

min

137

271

Improving Min and Max

- Suppose you have a priority queue where finding the min takes time $\omega(1)$.

- How could you modify it so that finding the min can be done in time $O(1)$?

- **Answer:** Store the minimum outside of the priority queue.

84

137

271

42

min

Improving Min and Max

- Suppose you have a priority queue where finding the min takes time ω(1).

- How could you modify it so that finding the min can be done in time O(1)?

- **Answer:** Store the minimum outside of the priority queue.
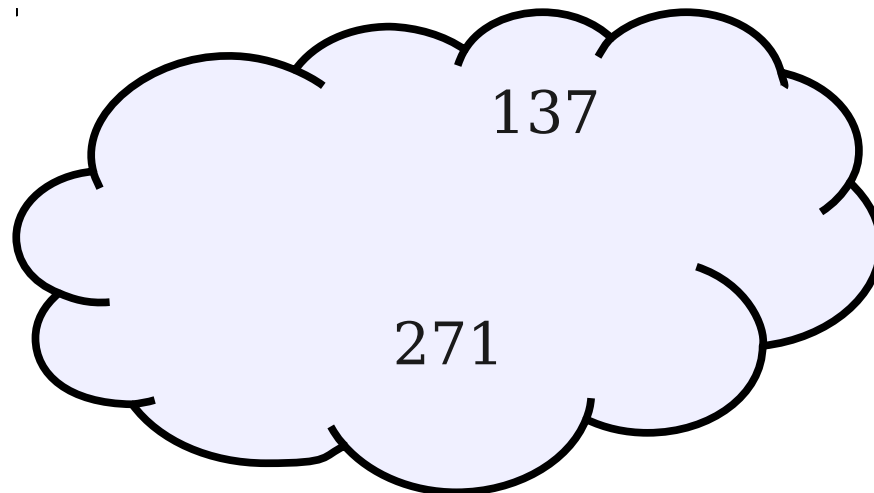
min

137

84

271

Improving Min and Max

- Suppose you have a priority queue where finding the min takes time $\omega(1)$.

- How could you modify it so that finding the min can be done in time $O(1)$?

- **Answer:** Store the minimum outside of the priority queue.
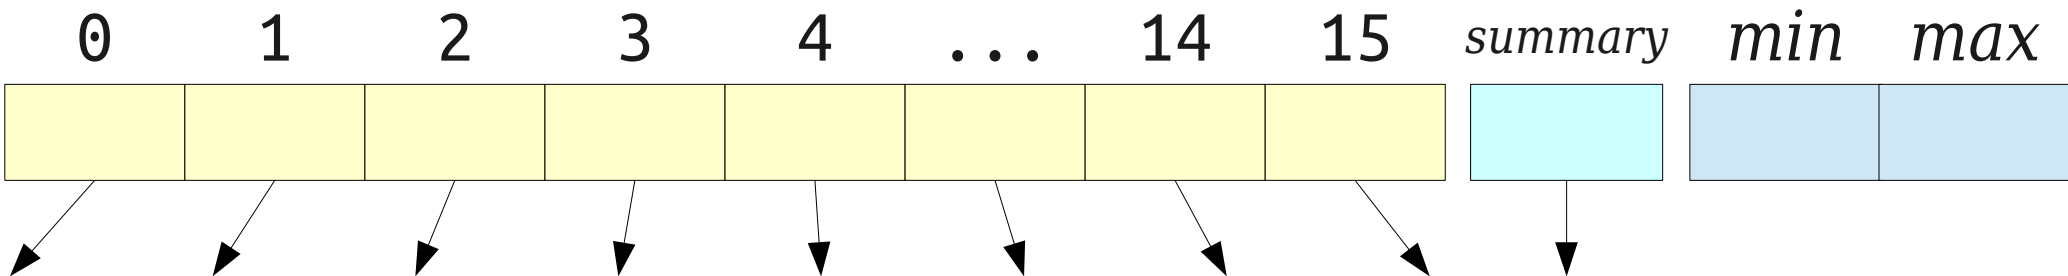
84

min

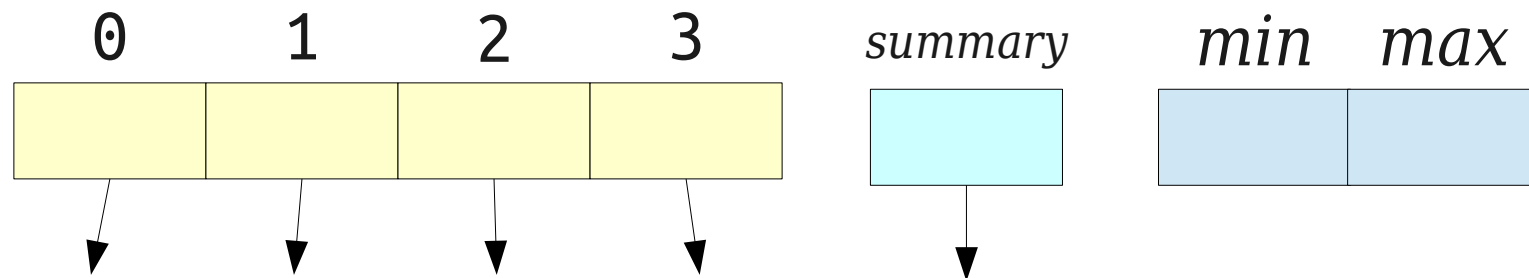137

271

van Emde Boas Trees

- A **van Emde Boas tree** is a slight modification to our previous structure.

- As before, split the universe into $\Theta(U^{1/2})$ blocks of size $\Theta(U^{1/2})$.

- As before, have the structure also store a summary of size $\Theta(U^{1/2})$.

- Additionally, have the data structure store the minimum and maximum separately from the rest of the structure.

- Each recursive copy of the data structure stores its own min and max storing the min and max value in its substructure.

van Emde Boas Trees

- Let $\mathcal{U} = [256]$.
- The top-level structure looks like this:

| 0 | 1 | 2 | 3 | 4 | ... | 14 | 15 | *summary* | *min* | *max* |

- Each structure one level below (and the summary) looks like this:

| 0 | 1 | 2 | 3 | *summary* | *min* | *max* |

vEB Tree Lookups

- Lookups in a vEB tree work as before, but with one extra step: check whether the value being searched for is the *min* or *max* value.

vEB Tree Insertions

- Insertions in a vEB tree work as before, but with extra logic to handle *min* and *max*.

  - May need to handle the case where the tree is empty.

vEB Tree Insertions

- Insertions in a vEB tree work as before, but with extra logic to handle *min* and *max*.

  - May need to handle the case where the tree is empty.

vEB Tree Insertions

- Insertions in a vEB tree work as before, but with extra logic to handle *min* and *max*.

  - May need to handle the case where the tree is empty.

  - May need to handle the case where the tree has just one element.

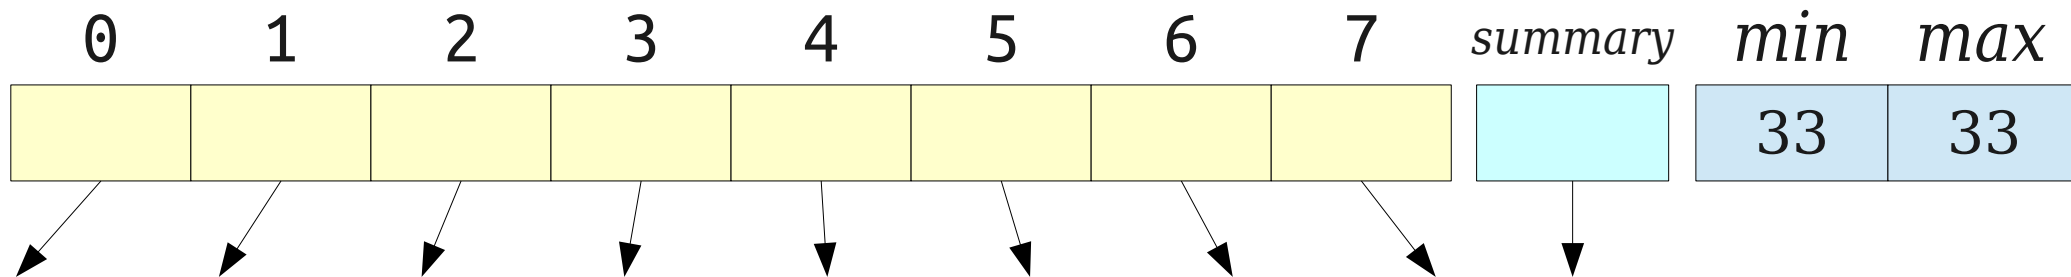| 0	1	2	3	4	5	6	7	*summary*	*min*	*max*
									33	33
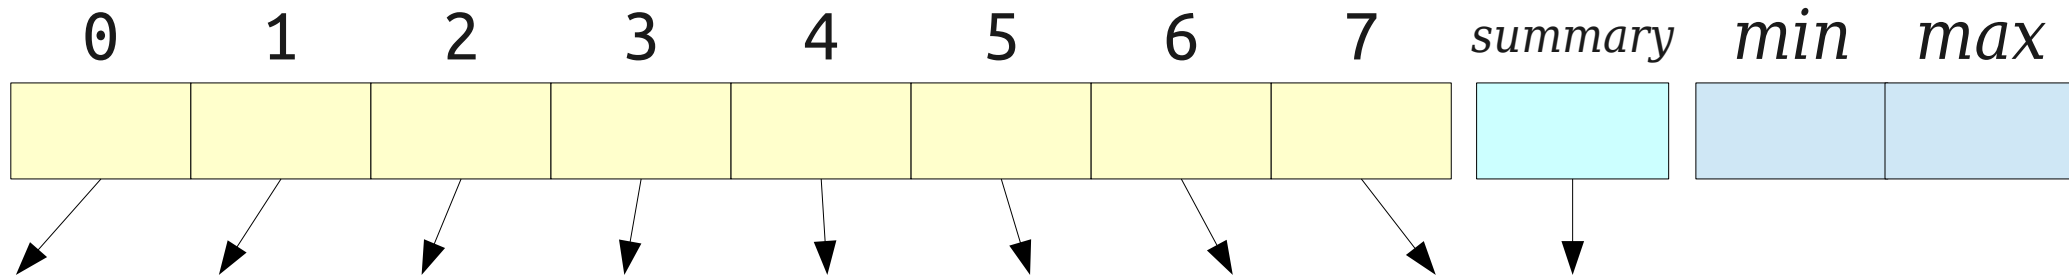
vEB Tree Insertions

- Insertions in a vEB tree work as before, but with extra logic to handle *min* and *max*.

  - May need to handle the case where the tree is empty.

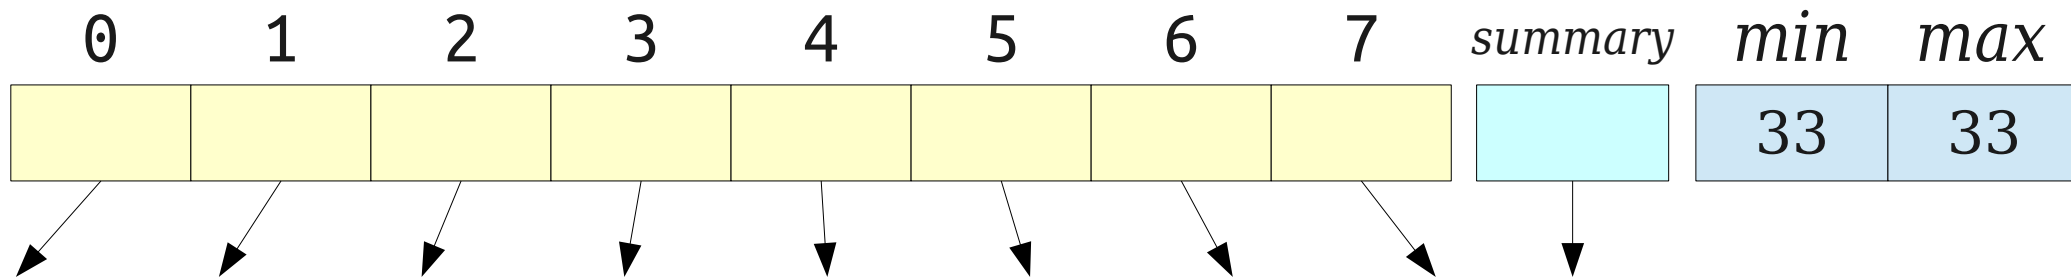 - May need to handle the case where the tree has just one element.

vEB Tree Insertions

- Insertions in a vEB tree work as before, but with extra logic to handle *min* and *max*.

  - May need to handle the case where the tree is empty.

  - May need to handle the case where the tree has just one element.

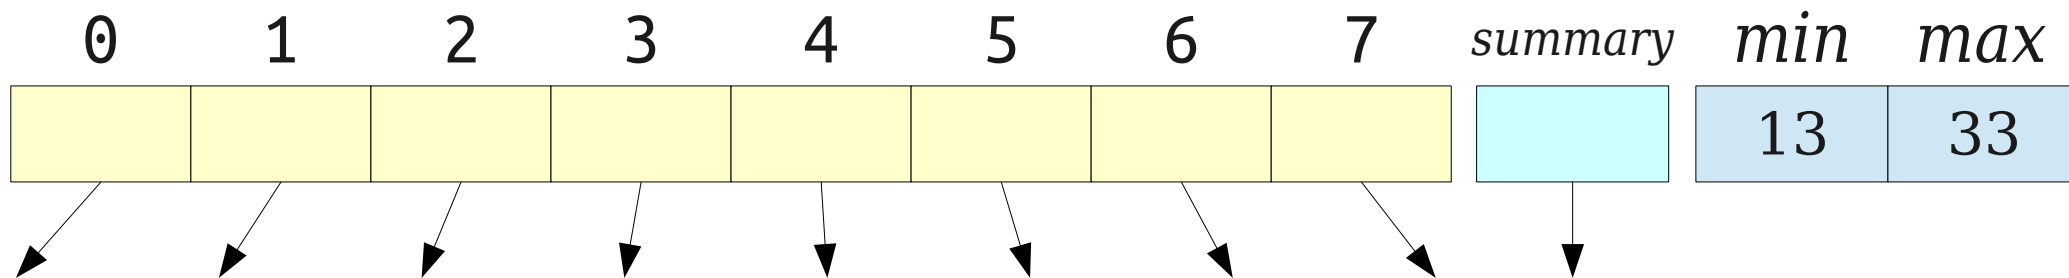 - May need to displace *min* or *max* into the tree.

vEB Tree Insertions

- Insertions in a vEB tree work as before, but with extra logic to handle *min* and *max*.
    - May need to handle the case where the tree is empty.
    - May need to handle the case where the tree has just one element.
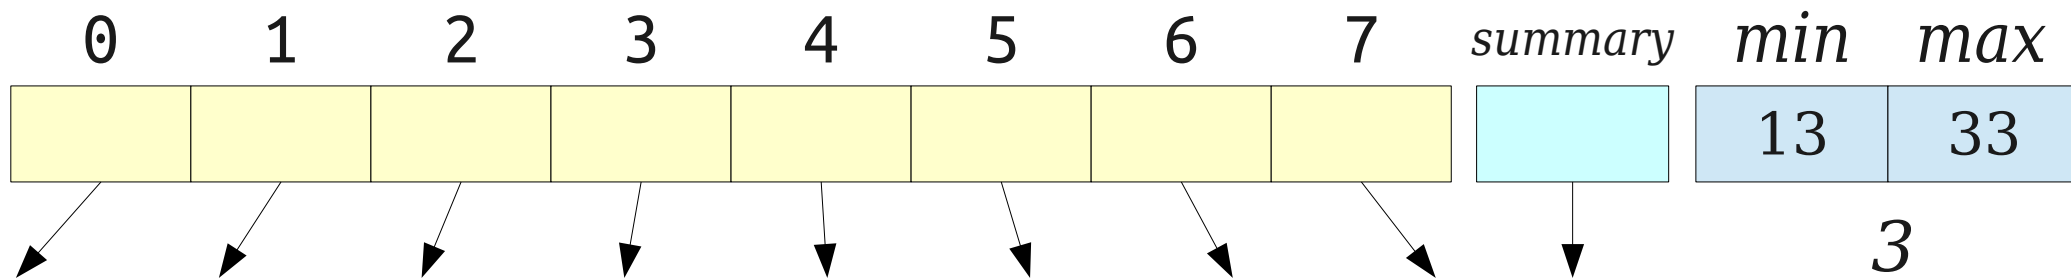   - May need to displace *min* or *max* into the tree.

vEB Tree Insertions

- Insertions in a vEB tree work as before, but with extra logic to handle *min* and *max*.

  - May need to handle the case where the tree is empty.

  - May need to handle the case where the tree has just one element.

  - May need to displace *min* or *max* into the tree.

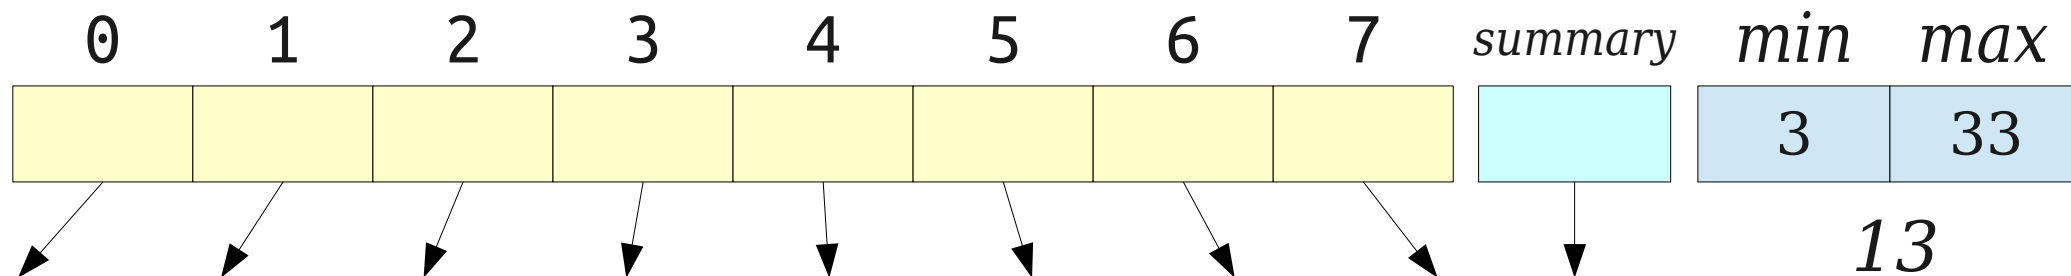| 0	1	2	3	4	5	6	7	*summary*	*min*	*max*
									3	33

13

vEB Tree Insertions

- Insertions in a vEB tree work as before, but with extra logic to handle *min* and *max*.

  - May need to handle the case where the tree is empty.

  - May need to handle the case where the tree has just one element.

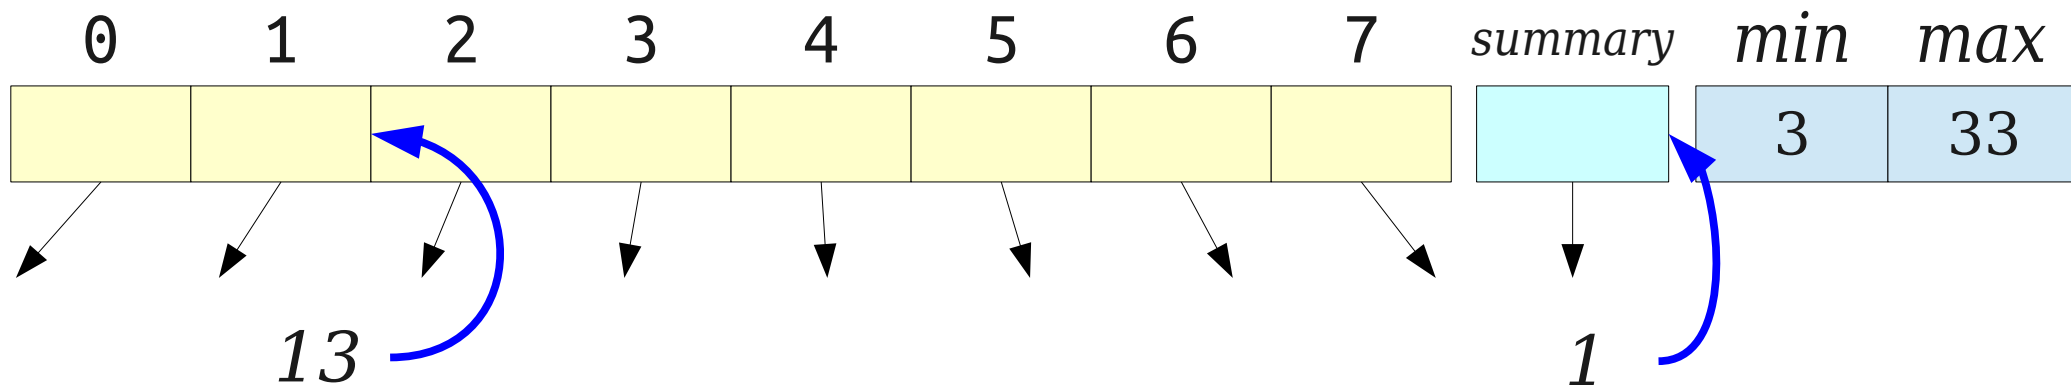 - May need to displace *min* or *max* into the tree.

vEB Tree Insertions

- Insertions in a vEB tree work as before, but with extra logic to handle *min* and *max*.

  - May need to handle the case where the tree is empty.

  - May need to handle the case where the tree has just one element.

  - May need to displace *min* or *max* into the tree.

vEB Tree Deletions

- Deletions in a vEB tree work as before, but with extra logic to handle *min* and *max*.

  - May need to pull an element to fill in a missing *min* or *max*.

0	1	2	3	4	5	6	7	*summary*	*min*	*max*
									3	33

vEB Tree Deletions

- Deletions in a vEB tree work as before, but with extra logic to handle *min* and *max*.
  - May need to pull an element to fill in a missing *min* or *max*.

vEB Tree Deletions

- Deletions in a vEB tree work as before, but with extra logic to handle *min* and *max*.

  - May need to pull an element to fill in a missing *min* or *max*.

We need to find the minimum element in these buckets.

0	1	2	3	4	5	6	7	*summary*	*min*	*max*
										33

vEB Tree Deletions

- Deletions in a vEB tree work as before, but with extra logic to handle *min* and *max*.

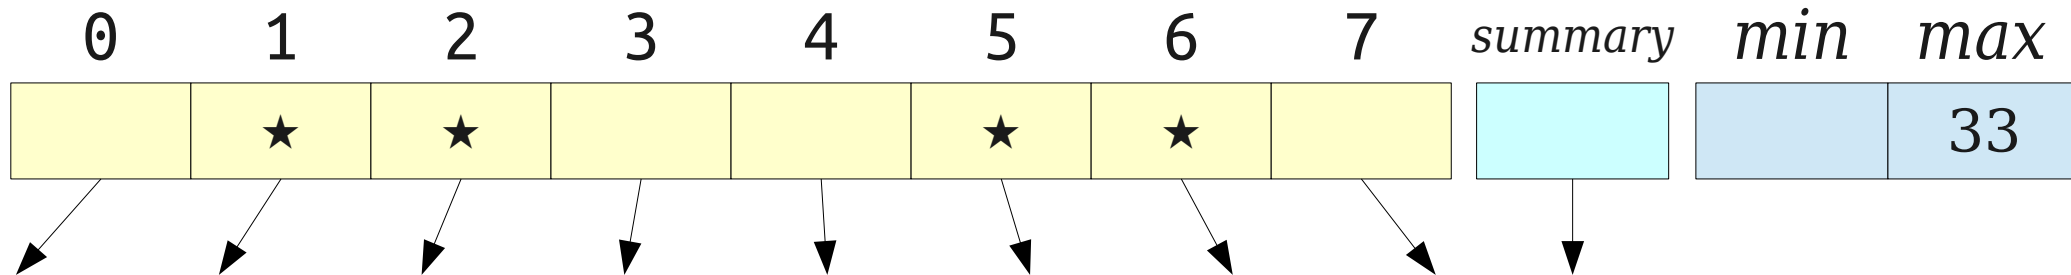 - May need to pull an element to fill in a missing *min* or *max*.

We need to find the minimum element in these buckets.

0	1	2	3	4	5	6	7	*summary*	*min*	*max*
	★	★			★	★				33

vEB Tree Deletions

- Deletions in a vEB tree work as before, but with extra logic to handle *min* and *max*.

  - May need to pull an element to fill in a missing *min* or *max*.

We need to find the minimum element in these buckets.

Ask the summary for the first nonempty block...

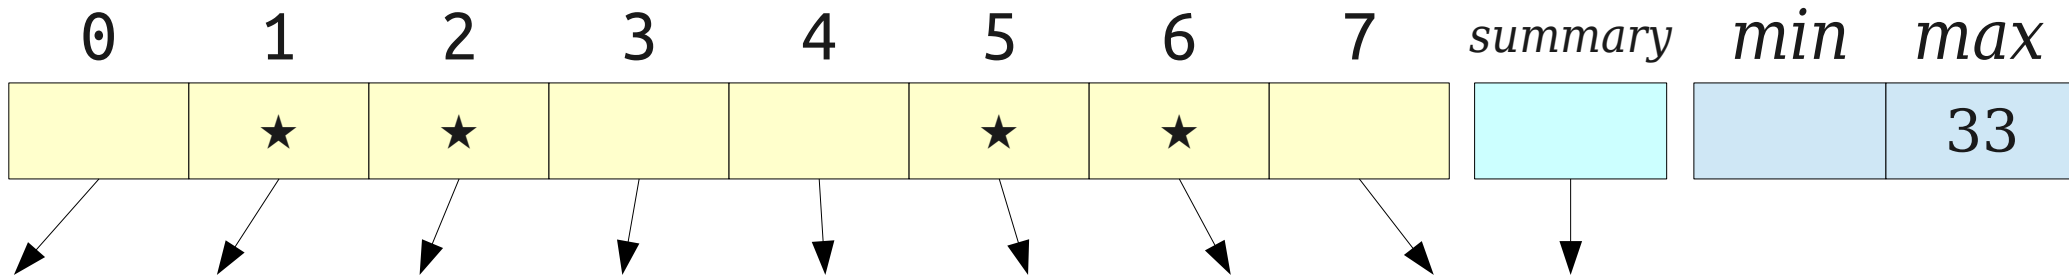| 0	1	2	3	4	5	6	7	summary	*min*	*max*
	★	★			★	★				33

vEB Tree Deletions

- Deletions in a vEB tree work as before, but with extra logic to handle *min* and *max*.

  - May need to pull an element to fill in a missing *min* or *max*.

We need to find the minimum element in these buckets.

Ask the summary for the first nonempty block...

0	1	2	3	4	5	6	7	*summary*	*min*	*max*
	★	★			★	★				33

...then delete its minimum and pull *min* up here.

vEB Tree Deletions

- Deletions in a vEB tree work as before, but with extra logic to handle *min* and *max*.

  - May need to pull an element to fill in a missing *min* or *max*.

We need to find the minimum element in these buckets.
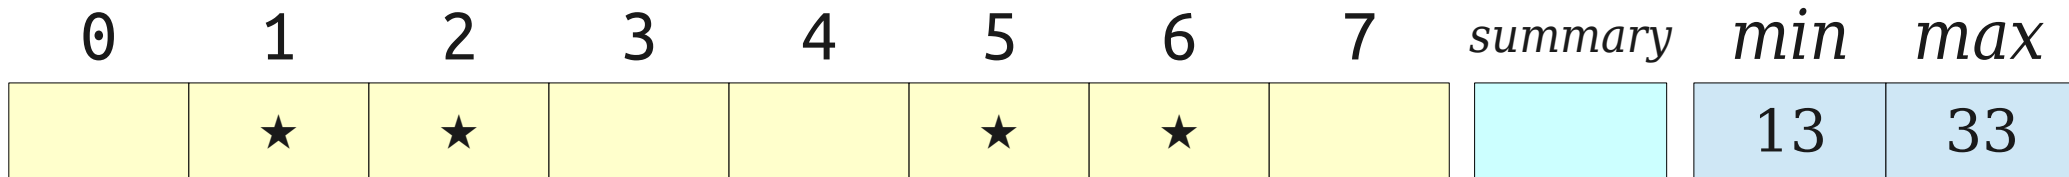
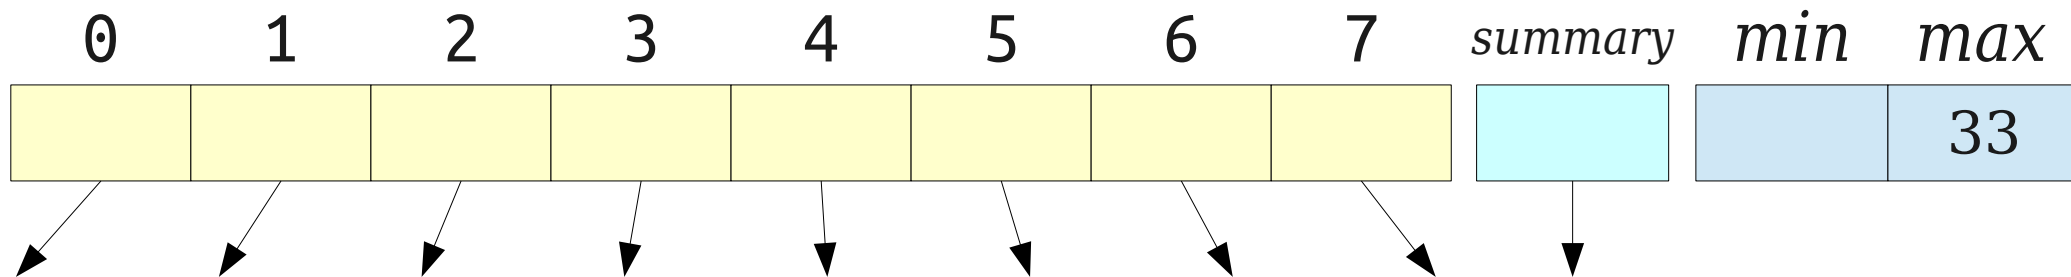Ask the summary for the first nonempty block...

0	1	2	3	4	5	6	7	*summary*	*min*	*max*
	★	★			★	★			13	33

...then delete its minimum and pull *min* up here.

vEB Tree Deletions

- Deletions in a vEB tree work as before, but with extra logic to handle *min* and *max*.
  - May need to pull an element to fill in a missing *min* or *max*.

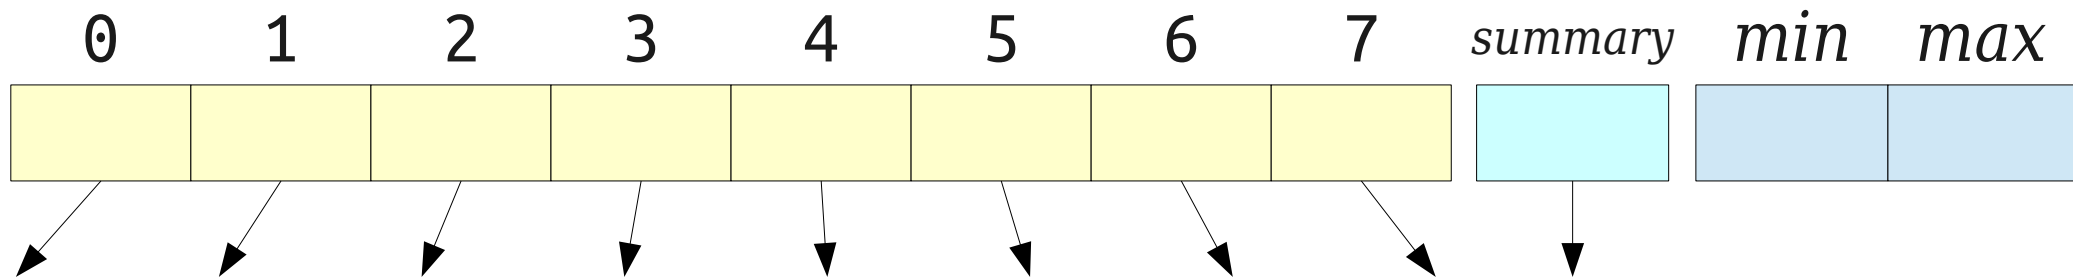| 0	1	2	3	4	5	6	7	*summary*	*min*	*max*
									13	33

vEB Tree Deletions

- Deletions in a vEB tree work as before, but with extra logic to handle *min* and *max*.

  - May need to pull an element to fill in a missing *min* or *max*.

  - May need to clear *min* or *max*.

vEB Tree Deletions

- Deletions in a vEB tree work as before, but with extra logic to handle *min* and *max*.

  - May need to pull an element to fill in a missing *min* or *max*.

  - May need to clear *min* or *max*.

vEB Tree Deletions

- Deletions in a vEB tree work as before, but with extra logic to handle *min* and *max*.

  - May need to pull an element to fill in a missing *min* or *max*.

  - May need to clear *min* or *max*.

vEB Tree Deletions

- Deletions in a vEB tree work as before, but with extra logic to handle *min* and *max*.

  - May need to pull an element to fill in a missing *min* or *max*.

  - May need to clear *min* or *max*.

Analyzing the Runtime

- This simple change profoundly affects the runtime of the operations for several reasons:
    - We can now instantly query for the min and max values in a tree.
    - The behavior of insert and delete changes slightly when working with empty or nearly empty trees.
- **min**, **max**, and **is-empty** run in time O(1).
- **lookup** runs in time O(log log U) as before.
- Let's revisit all the operations to see how efficiently they work.

Updating *insert*

- The logic for *insert*(x) works as follows:
  - If the tree is empty or has just one element, update *min* and *max* appropriately and stop.
  - Potentially displace the *min* or *max* and insert that value instead of x.
  - Insert $x \bmod U^{1/2}$ into the appropriate substructure.
  - Insert $\lfloor x / U^{1/2} \rfloor$ into the summary.
- Recurrence relation:

$$T(2) = \Theta(1)$$
$$T(U) = 2T(U^{1/2}) + \Theta(1).$$

- Still solves to O(log U). Can we do better?

An Observation

- The summary structure stores the indices of the substructures that are nonempty.

- Therefore, we only need to insert $\lfloor x / U^{1/2} \rfloor$ into the summary if that block previously was empty.

- Here's our new approach:

  - If the $\lfloor x / U^{1/2} \rfloor$th substructure is not empty:

    – Call **insert**$(x \bmod U^{1/2})$ into that substructure.

  - Otherwise:

    – Call **insert**$(x \bmod U^{1/2})$ into that substructure.

    – Call **insert**$(\lfloor x / U^{1/2} \rfloor)$ into the summary structure.

A Very Clever Insight

- **Useful Fact:** Inserting an element into an empty vEB tree takes time O(1).

A Very Clever Insight

- **Useful Fact:** Inserting an element into an empty vEB tree takes time O(1).
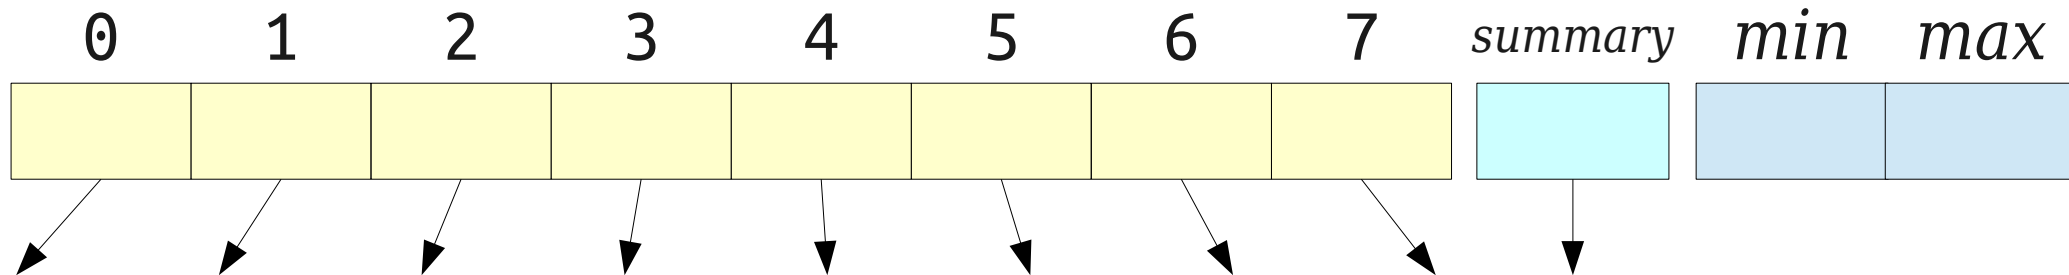
A Very Clever Insight

- **Useful Fact:** Inserting an element into an empty vEB tree takes time O(1).

- We only make at most one "real" recursive call:
  - If we don't recurse into the summary, we only made one recursive call down into a substructure.
  - If we make a recursive call into the summary, we did so because the other call was on an empty subtree, which isn't a "real" recursive call.

- New recurrence relation:

$$T(2) = \Theta(1)$$
$$T(U) \le T(U^{1/2}) + \Theta(1)$$

- As we've seen, this solves to **O(log log U)**.

Analyzing *delete*

- The logic for ***delete***(x) works as follows:
  - If the tree has just one element, update *min* and *max* appropriately and stop.
  - If *min* or *max* are being deleted, replace them with the *min* or *max* of the first or last nonempty tree, then proceed as if deleting that element instead.
  - Delete $x \bmod U^{1/2}$ from its subtree.
  - If that subtree is empty, delete $\lfloor x / U^{1/2} \rfloor$ from the summary.
- Recurrence relation:

$$T(2) = \Theta(1)$$
$$T(U) \leq 2T(U^{1/2}) + \Theta(1).$$

- Still solves to O(log U). However, is this bound tight?

A Better Analysis

- **Observation:** Deleting the last element out of a vEB tree takes time O(1).
  - Just need to update the *min* and *max* fields.
- Therefore, ***delete*** makes at most one "real" recursive call:
  - If it empties a subtree, the recursive call that did so ran in time O(1) and the "real" call is on the summary structure.
  - If it doesn't, then there's no second call on the summary structure.

The New Runtime

- With this factored in, the runtime of doing an **delete** is given by the recurrence

$$T(2) = \Theta(1)$$
$$T(U) \leq T(U^{1/2}) + \Theta(1)$$

- As we've seen, this solves to **O(log log U)**.

Finding a Successor

- In a vEB tree, we can find a successor as follows:
  - If the tree is empty or $x > max()$, there is no successor.
  - Otherwise, let i be the index of the tree containing x.
  - If subtree i is nonempty and x is less than i's max, x's successor is the successor in subtree i.
  - Otherwise, find the successor j of i in the summary.
  - If j exists, return the minimum value in tree j.
  - Otherwise, return the tree max.
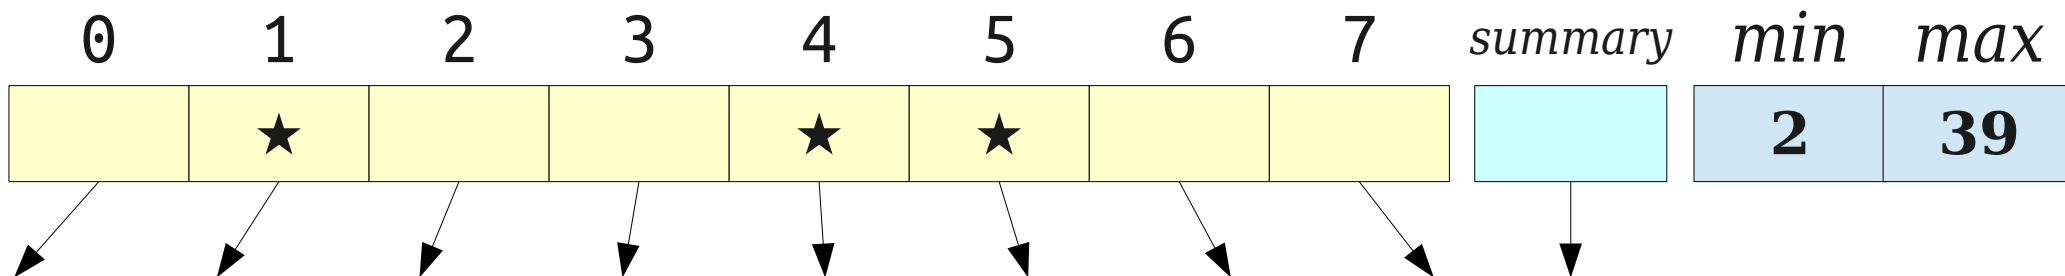
Finding a Successor

- In a vEB tree, we can find a successor as follows:
  - If the tree is empty or $x > \textbf{\textit{max}}()$, there is no successor.
  - Otherwise, let i be the index of the tree containing x.
  - If subtree i is nonempty and x is less than i's max, x's successor is the successor in subtree i.
  - Otherwise, find the successor j of i in the summary.
  - If j exists, return the minimum value in tree j.
  - Otherwise, return the tree max.
- At most one recursive call is made and each other operation needed runs in time $O(1)$.
- Recurrence: $T(U) \leq T(U^{1/2}) + \Theta(1)$; solves to **$O(\log \log U)$**.

van Emde Boas Trees

- The van Emde Boas tree supports insertions, deletions, lookups, successor queries, and predecessor queries in time O(log log U).

- It can answer min, max, and is-empty queries in time O(1).

- If $n = \omega(\log U)$, this is ***exponentially faster*** than a balanced BST!

The Catch

- There is, unfortunately, one way in which vEB trees stumble: **space usage**.

- We've assumed that the complete vEB tree has been constructed before we make any queries on it.

- How much space does it use?

The Recurrence

- The space usage of a van Emde Boas tree is given by the following recurrence relation:

$$S(2) = \Theta(1)$$

$$S(U) = (U^{1/2} + 1)S(U^{1/2}) + \Theta(U^{1/2})$$

- Using the substitution method, this can be shown to be $\Theta(U)$.

- Space usage is proportional to the size of the universe, not the number of elements stored!

Challenge:

Can we match the time bounds on van Emde Boas trees, but use $o(U)$ space?

Next Time

- **x-Fast Tries**

  - A randomized data structure matching the vEB bounds and using O($n \log U$) space.

- **y-Fast Tries**

  - A randomized data structure matching the vEB bounds in an amortized sense and using O(n) space.

- These data structures pull together just about everything we've covered this quarter – I hope they make for great midterm review!