

Range Minimum Queries


Part Two

Recap from Last Time

The RMQ Problem

- The **Range Minimum Query (RMQ)** problem is the following:

Given a fixed array A and two indices $i \leq j$, what is the smallest element out of $A[i], A[i + 1], \dots, A[j - 1], A[j]$?




Some Notation

- We'll say that an RMQ data structure has time complexity $\langle p(n), q(n) \rangle$ if
 - preprocessing takes time at most $p(n)$ and
 - queries take time at most $q(n)$.
- Last time, we saw structures with the following runtimes:
 - $\langle O(n^2), O(1) \rangle$ (full preprocessing)
 - $\langle O(n \log n), O(1) \rangle$ (sparse table)
 - $\langle O(n \log \log n), O(1) \rangle$ (hybrid approach)
 - $\langle O(n), O(n^{1/2}) \rangle$ (blocking)
 - $\langle O(n), O(\log n) \rangle$ (hybrid approach)
 - $\langle O(n), O(\log \log n) \rangle$ (hybrid approach)

The Framework


- Split the input into blocks of size b .
- Form an array of the block minima.
- Construct a “summary” RMQ structure over the block minima.
- Construct “block” RMQ structures for each block.
- Aggregate the results together.



The Framework

- Suppose we use a $\langle p_1(n), q_1(n) \rangle$ -time RMQ solution for the block minima and a $\langle p_2(n), q_2(n) \rangle$ -time RMQ solution within each block. Let the block size be b .
- In the hybrid structure, the preprocessing time is

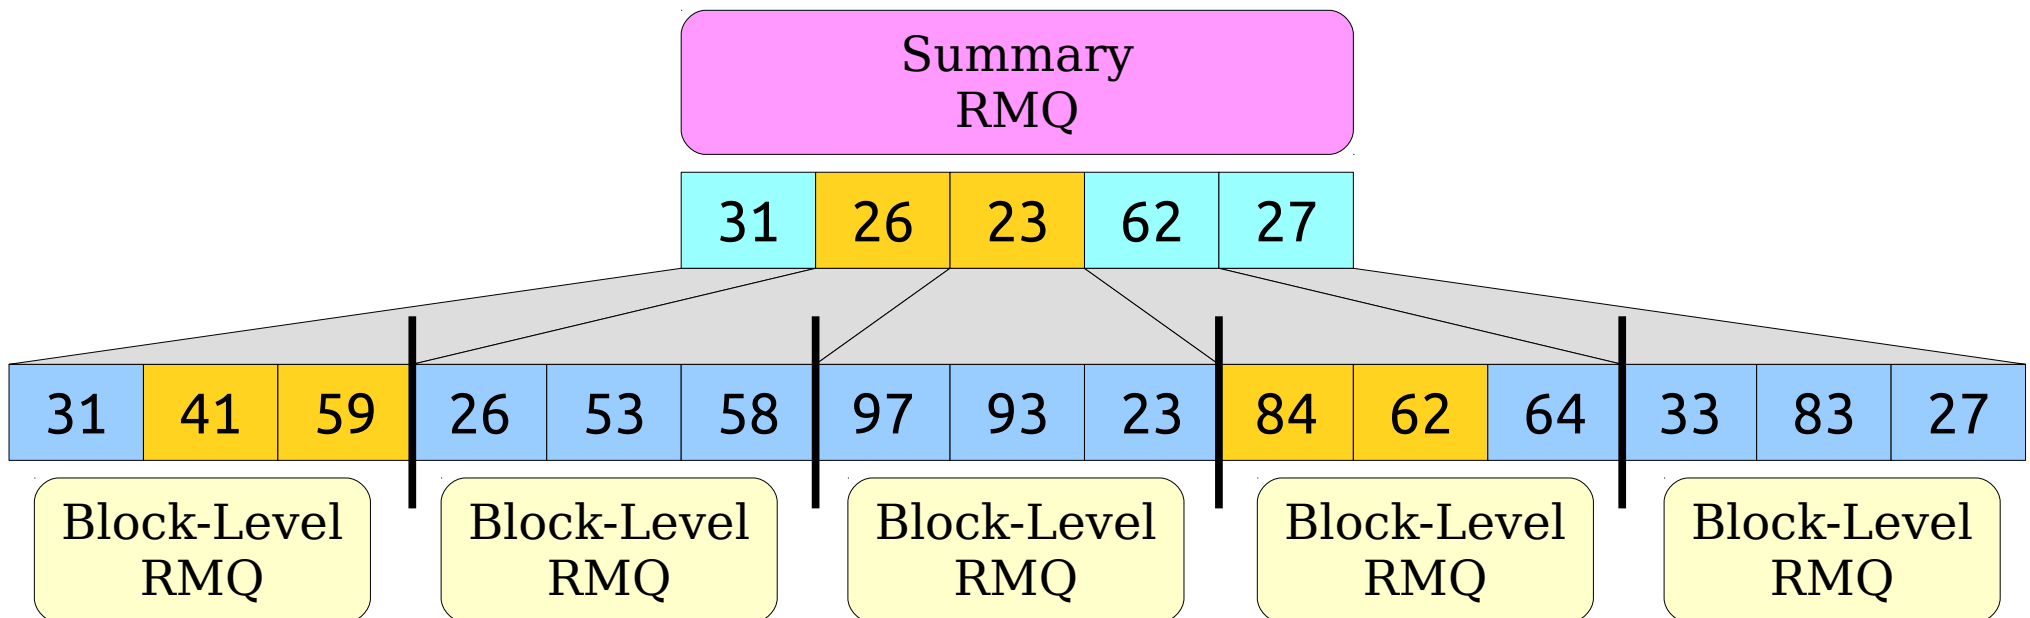
$$O(n + p_1(n / b) + (n / b) p_2(b))$$



The Framework

- Suppose we use a $\langle p_1(n), q_1(n) \rangle$ -time RMQ solution for the block minima and a $\langle p_2(n), q_2(n) \rangle$ -time RMQ solution within each block. Let the block size be b .
- In the hybrid structure, the query time is

$$O(q_1(n / b) + q_2(b))$$



Is there an $\langle O(n), O(1) \rangle$ solution to RMQ?

Yes!

New Stuff!

An Observation

The Limits of Hybrids

- The preprocessing time on a hybrid structure is

$$O(n + p_1(n / b) + (n / b) p_2(b)).$$

- The query time is

$$O(q_1(n / b) + q_2(b)).$$

- To build an $\langle O(n), O(1) \rangle$ hybrid, we need to have $p_2(n) = O(n)$ and $q_2(n) = O(1)$.

- ***We can't build an optimal solution with the hybrid approach unless we already have one!***

- ***Or can we?***

A Key Difference

- Our original problem is

Solve RMQ on a single array in time $\langle O(n), O(1) \rangle$

- The new problem is

Solve RMQ on a large number of small arrays with $O(1)$ query time and *total* preprocessing time $O(n)$.

- These are not the same problem.
- **Question:** Why is this second problem any easier than the first?

An Observation

10	30	20	40
----	----	----	----

166	361	261	464
-----	-----	-----	-----

Claim: The indices of the answers to any range minimum queries on these two arrays are the same.

Modifying RMQ

- From this point forward, let's have $\text{RMQ}_A(i, j)$ denote the **index** of the minimum value in the range rather than the value itself.
- **Observation:** If RMQ structures return indices rather than values, we can use a single RMQ structure for both of these arrays:

10	30	20	40
----	----	----	----

166	361	261	464
-----	-----	-----	-----

Some Notation

- Let B_1 and B_2 be blocks of length b .
- We'll say that B_1 and B_2 **have the same block type** (denoted $B_1 \sim B_2$) if the following holds:

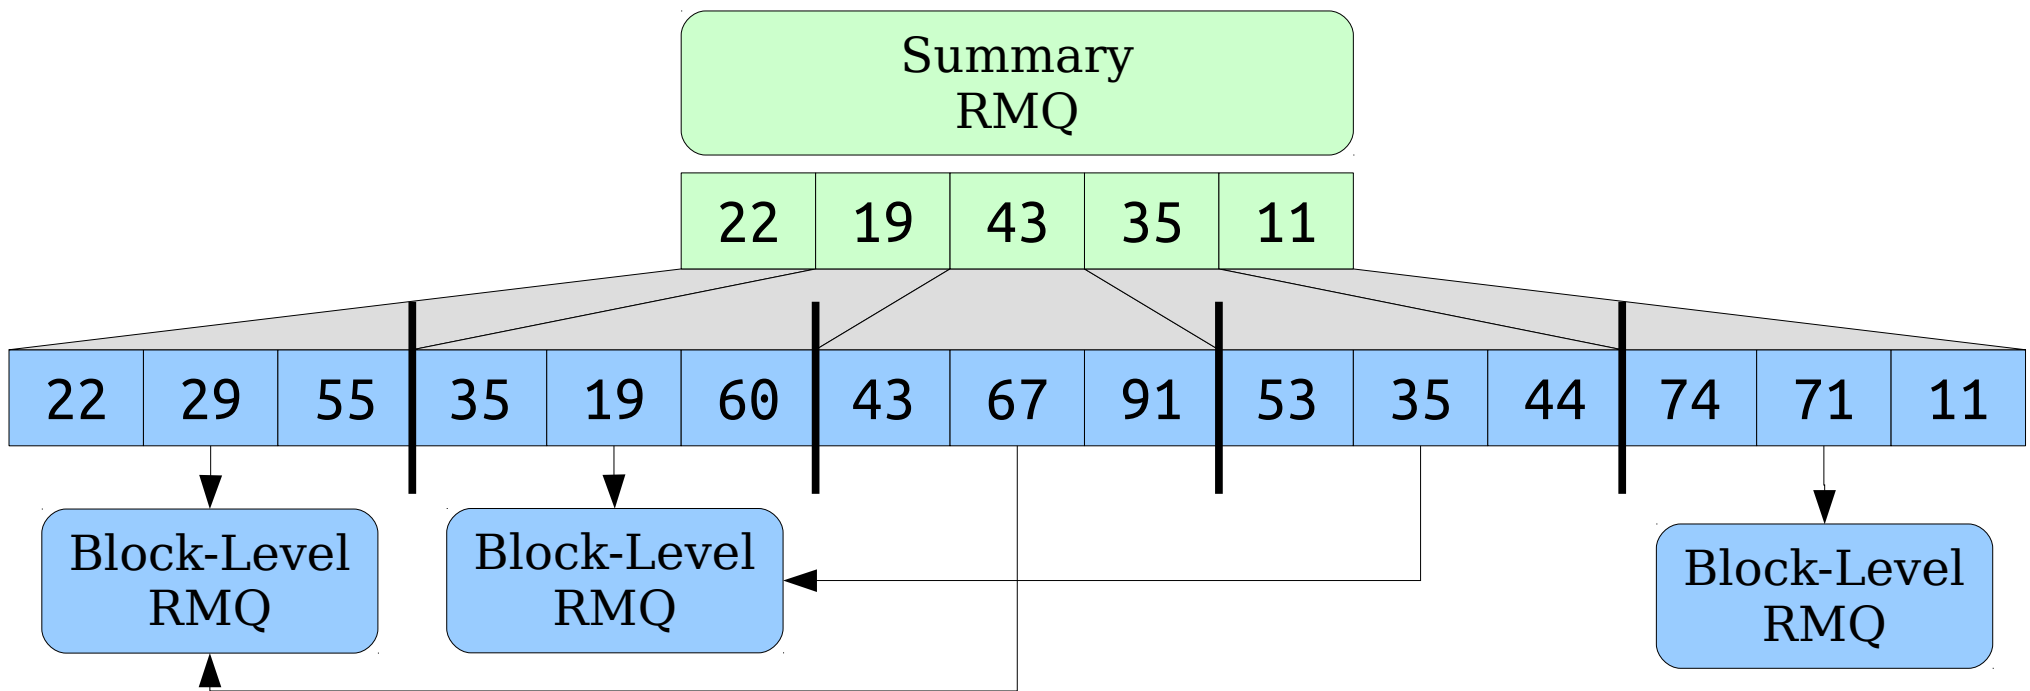
$$\text{For all } 0 \leq i \leq j < b: \\ \text{RMQ}_{B_1}(i, j) = \text{RMQ}_{B_2}(i, j)$$

- Intuitively, the RMQ answers for B_1 are always the same as the RMQ answers for B_2 .
- If we build an RMQ to answer queries on some block B_1 , we can reuse that RMQ structure on some other block B_2 iff $B_1 \sim B_2$.

Where We're Going

- Suppose we use an $\langle O(n \log n), O(1) \rangle$ sparse table for the top and the $\langle O(n^2), O(1) \rangle$ precompute-all structures for the blocks.
- However, whenever possible, we share block-level RMQ structures across multiple blocks.
- Our preprocessing time is now

$O(n + (n / b) \log n + b^2 \cdot \#distinct\text{-blocks-of-size-}b)$



Two Big Questions

***How can we tell when two blocks
can share RMQ structures?***

(Without an answer, this whole approach doesn't work!)

***How many block types are there,
as a function of b ?***

(Without an answer, we can't measure efficiency!)

The Adventure Begins!

Detecting Block Types

- For this approach to work, we need to be able to check whether two blocks have the same block type.
- **Problem:** Our formal definition of $B_1 \sim B_2$ is defined in terms of RMQ.
 - Not particularly useful *a priori*; we don't want to have to compute RMQ structures on B_1 and B_2 to decide whether they have the same block type!
- Is there a simpler way to determine whether two blocks have the same type?

An Initial Idea

- Since the elements of the array are ordered and we're looking for the smallest value in certain ranges, we might look at the permutation types of the blocks.

31	41	59	16	18	3	27	18	28	66	73	84
1	2	3	2	3	1	2	1	3	1	2	3
12	2	5	66	26	6	60	22	14	72	99	27
3	1	2	3	2	1	3	2	1	2	3	1

- **Claim:** If B_1 and B_2 have the same permutation on their elements, then $B_1 \sim B_2$.

Some Problems

- There are two main problems with this approach.
- **Problem One:** It's possible for two blocks to have different permutations but the same block type.
- All three of these blocks have the same block type but different permutation types:

261	268	161	167	166	167	261	161	268	166	166	268	161	261	167
4	5	1	3	2	3	4	1	5	2	2	5	1	4	3

- **Problem Two:** The number of possible permutations of a block is $b!$.
 - b has to be absolutely minuscule for $b!$ to be small.
- Is there a better criterion we can use?

An Observation

- **Claim:** If $B_1 \sim B_2$, the minimum elements of B_1 and B_2 must occur at the same position.

261	268	161	167	166
-----	-----	-----	-----	-----




75	35	80	85	83
----	----	----	----	----

6	5	3	9	7
---	---	---	---	---

14	22	11	43	35
----	----	----	----	----

An Observation


- **Claim:** If $B_1 \sim B_2$, the minimum elements of B_1 and B_2 must occur at the same position.



- **Claim:** This property must hold recursively on the subarrays to the left and right of the minimum.

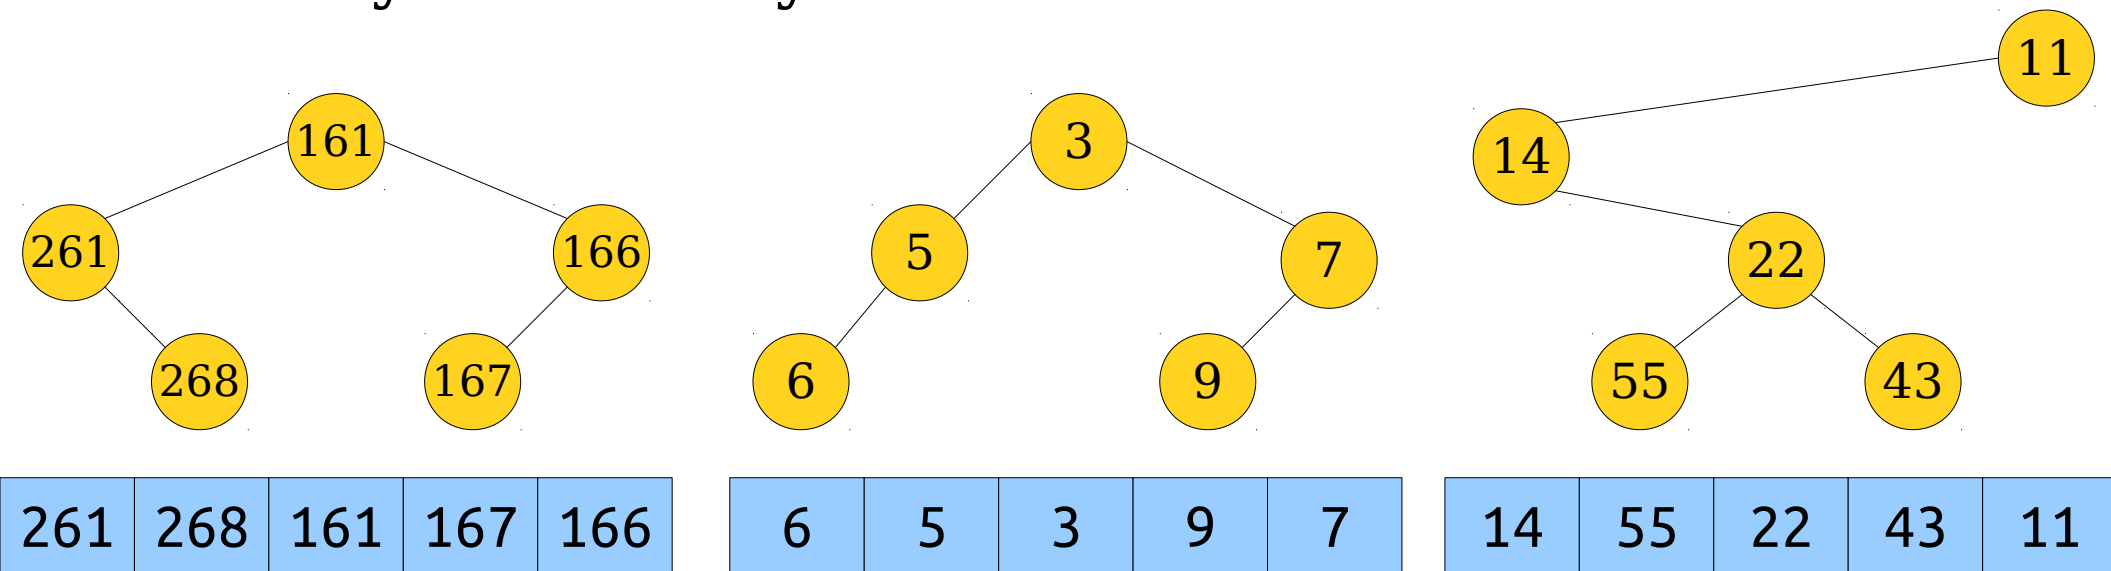
Cartesian Trees

- A **Cartesian tree** is a binary tree derived from an array and defined as follows:
 - The empty array has an empty Cartesian tree.
 - For a nonempty array, the root stores the minimum value. Its left and right children are Cartesian trees for the subarrays to the left and right of the minimum.





Cartesian Trees

- A **Cartesian tree** can also be defined as follows:
 - The Cartesian tree for an array is a binary tree obeying the **min-heap property** whose inorder traversal gives back the original array.
- This holistic description of the tree will come in handy later today.





Cartesian Trees and RMQ

- **Theorem:** Let B_1 and B_2 be blocks of length b . Then $B_1 \sim B_2$ iff B_1 and B_2 have isomorphic Cartesian trees.




Cartesian Trees and RMQ

- **Theorem:** Let B_1 and B_2 be blocks of length b . Then $B_1 \sim B_2$ iff B_1 and B_2 have isomorphic Cartesian trees.
- **Proof sketch:**
 - (\Rightarrow) Induction. B_1 and B_2 have equal RMQs, so corresponding ranges have minima at the same positions.



Cartesian Trees and RMQ

- **Theorem:** Let B_1 and B_2 be blocks of length b . Then $B_1 \sim B_2$ iff B_1 and B_2 have isomorphic Cartesian trees.
- **Proof sketch:**
 - (\Leftarrow) Induction. It's possible to answer RMQ using a recursive walk on the Cartesian tree.



31	41	59	26	53	58	97	23	93	84	33	64	62	83	27
----	----	----	----	----	----	----	----	----	----	----	----	----	----	----

***How can we tell when two blocks
can share RMQ structures?***

***When those blocks have isomorphic Cartesian trees!
But how do we check that?***

***How many block types are there,
as a function of b ?***

~_(ツ)_/

How quickly can we build a Cartesian tree?

Building Cartesian Trees


- Here's a naïve algorithm for constructing Cartesian trees:
 - Find the minimum value.
 - Recursively build a Cartesian tree for the array to the left of the minimum.
 - Recursively build a Cartesian tree with the elements to the right of the minimum.
 - Return the overall tree.
- How efficient is this approach?

Building Cartesian Trees

- This algorithm works by
 - doing a linear scan over the array,
 - identifying the minimum at whatever position it occupies, then
 - recursively processing the left and right halves on the array.
- Similar to the recursion in quicksort: it depends on where the minima are.
 - Always get good splits: $\Theta(n \log n)$.
 - Always get bad splits: $\Theta(n^2)$.
- We're going to need to be faster than this.

A Better Approach


- It turns out that it's possible to build a Cartesian tree over an array of length k in time $O(k)$.
- **High-level idea:** Build a Cartesian tree for the first element, then the first two, then the first three, then the first four, etc.



Observation 1: This new node must end up as the rightmost node on the right spine. (An inorder traversal of the Cartesian tree gives the original array.)

A Better Approach


- It turns out that it's possible to build a Cartesian tree over an array of length k in time $O(k)$.
- **High-level idea:** Build a Cartesian tree for the first element, then the first two, then the first three, then the first four, etc.



Observation 2: Cartesian trees are min-heaps (each node's value is at least as large as its parent's).

A Better Approach


- It turns out that it's possible to build a Cartesian tree over an array of length k in time $O(k)$.
- **High-level idea:** Build a Cartesian tree for the first element, then the first two, then the first three, then the first four, etc.



Idea: Reshape the right spine of the tree to put the new node into the right place.


A Better Approach

- It turns out that it's possible to build a Cartesian tree over an array of length k in time $O(k)$.
- **High-level idea:** Build a Cartesian tree for the first element, then the first two, then the first three, then the first four, etc.



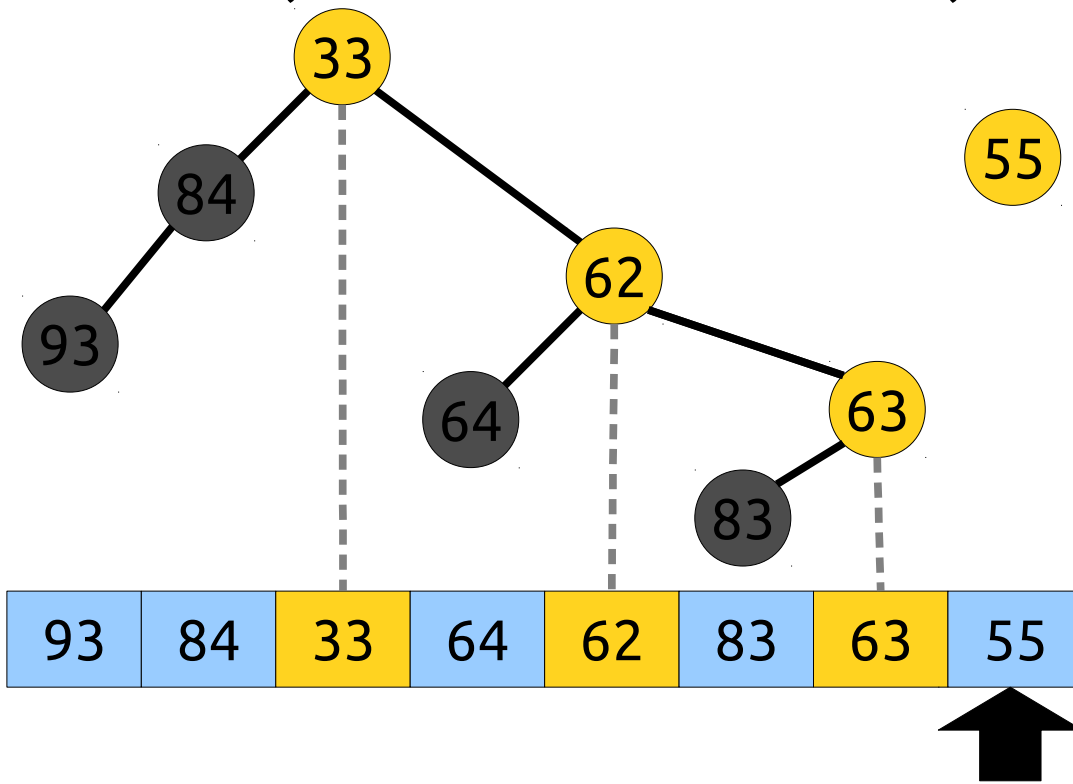
A Better Approach

- It turns out that it's possible to build a Cartesian tree over an array of length k in time $O(k)$.
- **High-level idea:** Build a Cartesian tree for the first element, then the first two, then the first three, then the first four, etc.




A Better Approach

- It turns out that it's possible to build a Cartesian tree over an array of length k in time $O(k)$.
- **High-level idea:** Build a Cartesian tree for the first element, then the first two, then the first three, then the first four, etc.



A Better Approach

- It turns out that it's possible to build a Cartesian tree over an array of length k in time $O(k)$.
- **High-level idea:** Build a Cartesian tree for the first element, then the first two, then the first three, then the first four, etc.



An Efficient Implementation

A Stack-Based Algorithm

- Maintain a stack of the nodes on the right spine of the tree.
- To insert a new node:
 - Pop the stack until it's empty or the top node has a lower value than the current value.
 - Set the new node's left child to be the last value popped (or **null** if nothing was popped).
 - Set the new node's parent to be the top node on the stack (or **null** if the stack is empty).
 - Push the new node onto the stack.



32	45	16	18	9	33
----	----	----	----	---	----

Analyzing the Runtime

- Adding in another node to the Cartesian tree might take time $O(n)$, since we might have to pop everything off the stack.
- Since there are n nodes to add, the runtime of this approach is $O(n^2)$.
- **Claim:** This is a weak bound! The runtime is actually $\Theta(n)$.
- **Proof:** Work done per node is proportional to the number of stack operations performed when that node was processed.
- Total number of stack operations is at most $2n$.
 - Every node is pushed once.
 - Every node is popped at most once.
- Total runtime is therefore $\Theta(n)$.

Time-Out for Announcements!

Office Hours

- We'll be holding four sets of office hours each week!
- ***Mondays:***
 - Anton's office hours: 10:00AM – 12:00PM, Huang Basement.
 - Michael's office hours: 1:30PM – 3:30PM, Huang Basement.
- ***Wednesdays:***
 - Keith's office hours: 2:00PM – 4:00PM, Gates 178.
- ***Fridays:***
 - Ryan's office hours: 1:00PM – 3:00PM, Huang Basement.

Problem Set Zero

- Problem Set Zero is due next Tuesday at 2:30PM.
- Have questions?
 - Ask on Piazza!
 - Stop by office hours!

Problem Set Logistics

- We will be using GradeScope for assignment submissions this quarter.
- To use it, visit the GradeScope website and use the code

M42XBN

to register for CS166.

- ***No hardcopy assignments will be accepted.*** We're using GradeScope to track due dates and as a gradebook.
- Programming assignments are submitted separately using our AFS-hosted submission script.

Back to CS166!

***How can we tell when two blocks
can share RMQ structures?***

***When those blocks have isomorphic Cartesian trees!
And we can check this in time $O(b)$!***

***How many block types are there,
as a function of b ?***

(ツ)/

Theorem: The number of Cartesian trees for an array of length b is at most 4^b .

In case you're curious, the actual number is

$$\frac{1}{b+1} \binom{2b}{b},$$

which is roughly equal to

$$\frac{4^b}{b^{3/2} \sqrt{\pi}}.$$

Look up the **Catalan numbers** for more information!


Proof Approach

- Our stack-based algorithm for generating Cartesian trees produces a Cartesian tree for every possible input array.
- Therefore, if we can count the number of possible executions of that algorithm, we can count the number of Cartesian trees.
- Using a simple counting scheme, we can show that there are at most 4^b possible executions.

The Insight

- **Claim:** The Cartesian tree produced by the stack-based algorithm is uniquely determined by the sequence of pushes and pops made on the stack.
- There are at most $2b$ stack operations during the execution of the algorithm: b pushes and no more than b pops.
- Can represent the execution as a $2b$ -bit number, where 1 means “push” and 0 means “pop.” We'll pad the end with 0's (pretend we pop everything from the stack.)
 - We'll call this number the **Cartesian tree number** of a particular block.
- There are at most $2^{2b} = 4^b$ possible $2b$ -bit numbers, so there are at most 4^b possible Cartesian trees.

Cartesian Tree Numbers



32	45	16	18	9	33
----	----	----	----	---	----

1	1	0	0	1	1	0	0	1	1	0	0
---	---	---	---	---	---	---	---	---	---	---	---

0 means pop; 1 means push

One Last Observation

- **Recall:** Our goal is to be able to detect when two blocks have the same type so that we can share RMQ structures between them.
- We've seen that two blocks have the same type if and only if they have the same Cartesian tree.
- Using the connection between Cartesian trees and Cartesian tree numbers, we can see that ***we don't actually have to build any Cartesian trees!***
- We can just compute the Cartesian tree number of each block and use those numbers to test for block equivalence.

Cartesian Tree Numbers

27	18	28	18	28	45	90	45	23	53	60	28	74	71	35
----	----	----	----	----	----	----	----	----	----	----	----	----	----	----

|

1 0 1 1 0 1 1 1 1 0 1 0 0 0 1 1 1 0 0 1 1 0 1 0 1 0 0 0 0

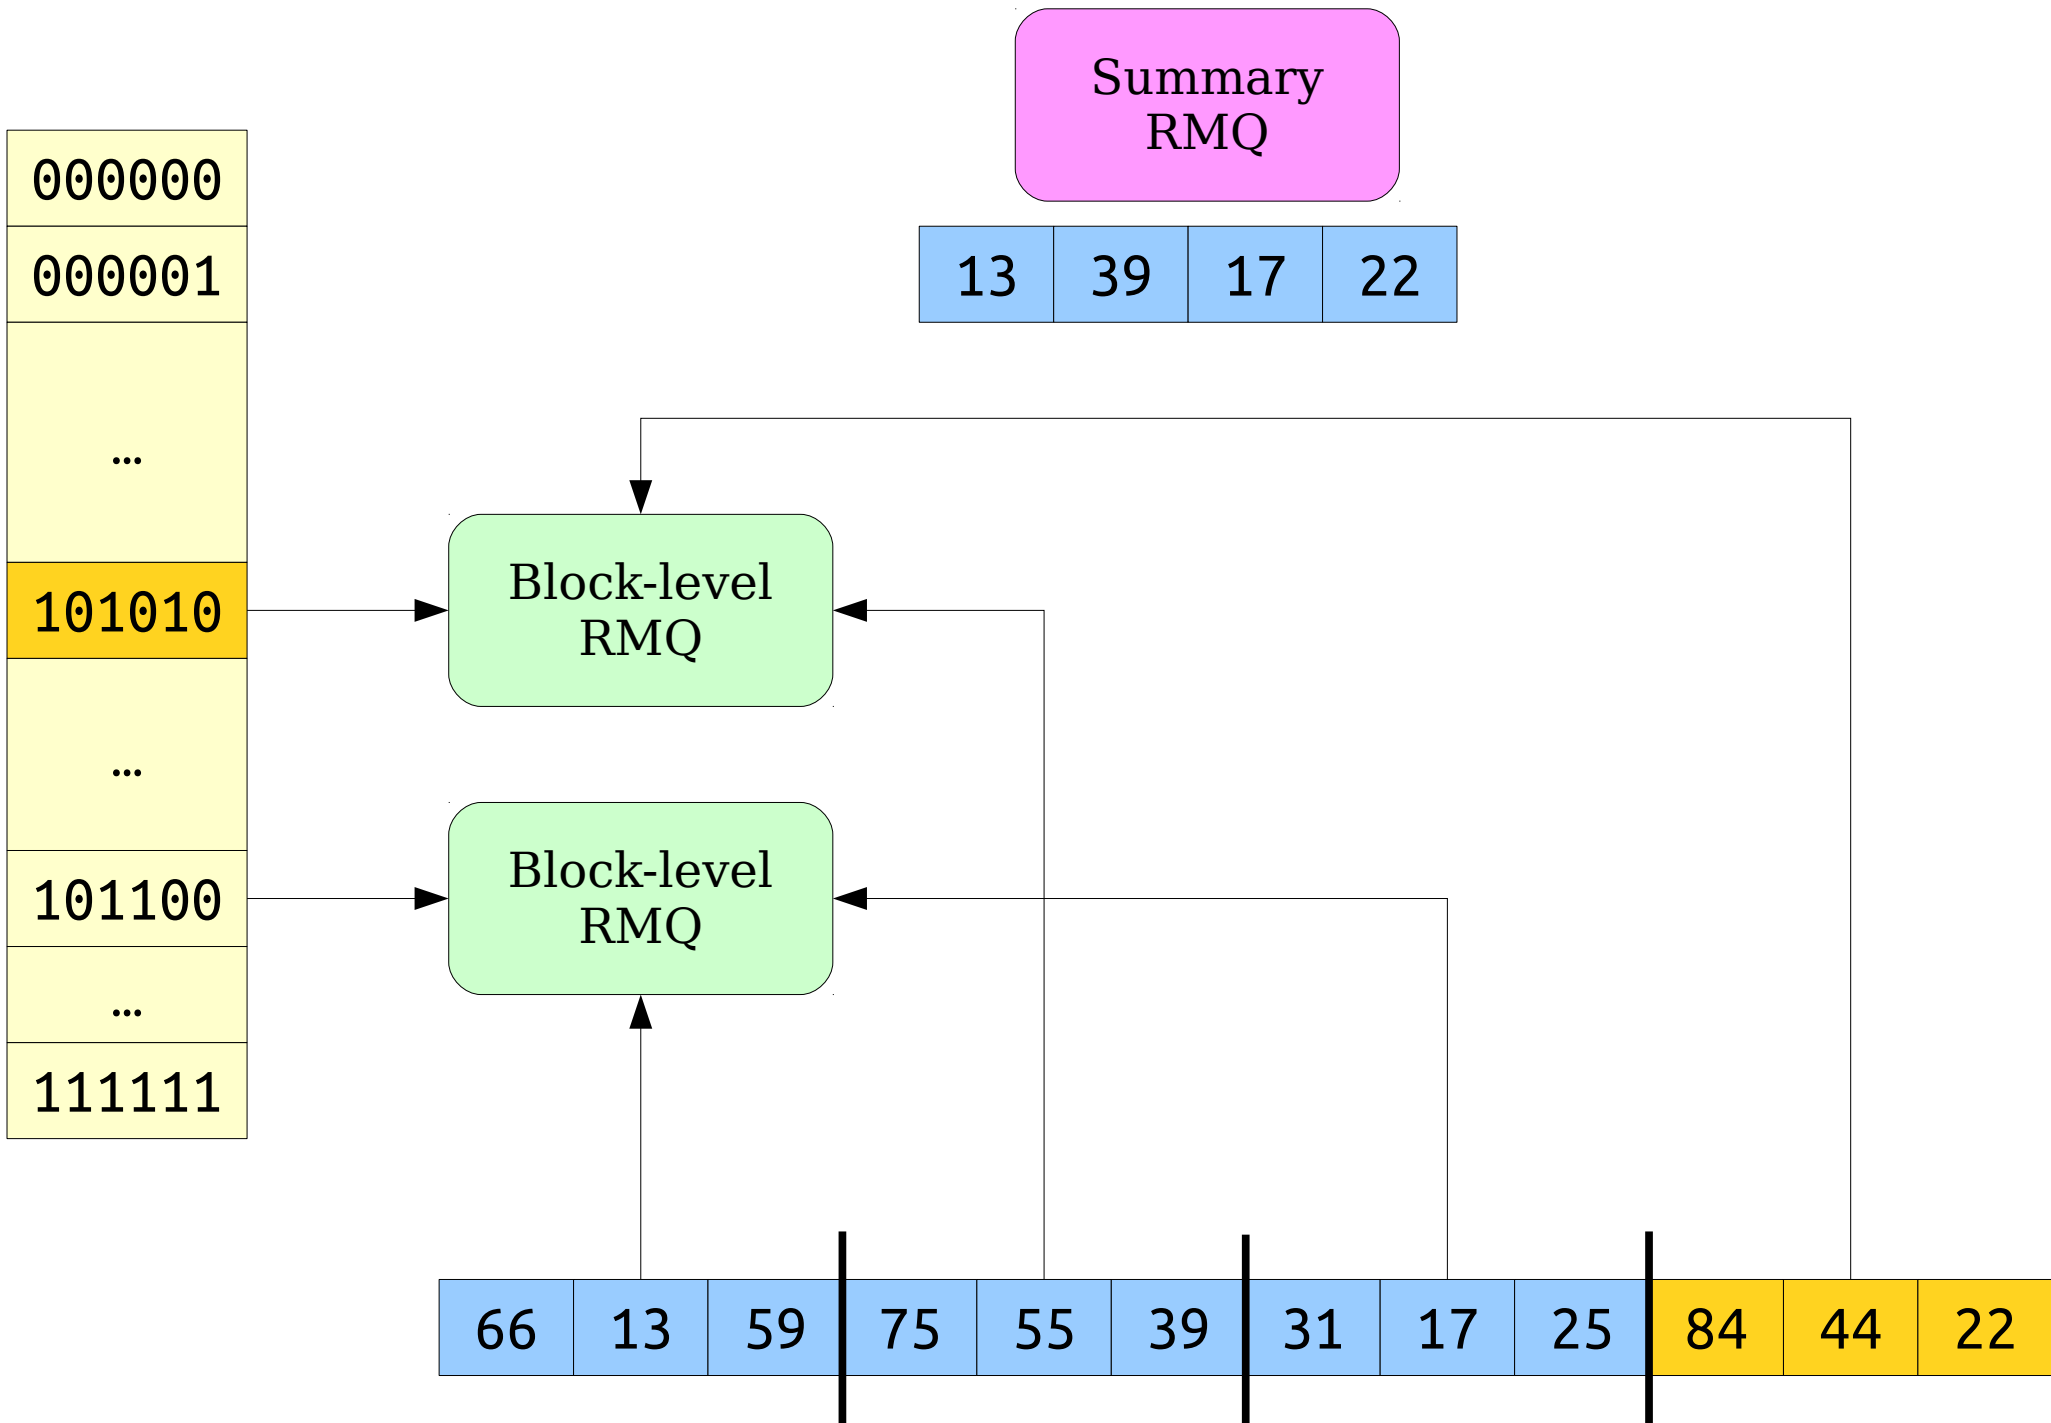
***How can we tell when two blocks
can share RMQ structures?***

***When they have the same Cartesian tree number!
And we can check this in time $O(b)$!***

***How many block types are there,
as a function of b ?***

At most 4^b , because of the above algorithm!

Putting it all Together



How Efficient is This?

We're using the hybrid approach,
and all the types we're using have
constant query times.

Query time: **$O(1)$**

The 4^b term grows exponentially in n unless we pick $b = O(\log n)$.

Our preprocessing time is

$$O(n + (n / b) \log n + b^2 4^b)$$

The $(n / b) \log n$ term will be superlinear unless we pick $b = \Omega(\log n)$.

Our preprocessing time is

$$O(n + (n / b) \log n + b^2 4^b)$$

Suppose we pick
 $b = k \log_4 n$
for some constant k .

Our preprocessing time is

$$O(n + n + b^2 4^b)$$

Suppose we pick
 $b = k \log_4 n$
for some constant k .

Our preprocessing time is

$$O(n + n + (\log n)^2 n^k)$$

Suppose we pick
 $b = k \log_4 n$
for some constant k .

Our preprocessing time is

$$O(n + n + (\log n)^2 n^k)$$

Suppose we pick
 $b = k \log_4 n$
for some constant k .

Now, set $k = 1/2$.

Our preprocessing time is

$$O(n)$$

Suppose we pick
 $b = k \log_4 n$
for some constant k .

Now, set $k = 1/2$.

The Fischer-Heun Structure

- This approach is called the ***Fischer-Heun structure***. It uses a modified version of our hybrid approach:
 - Set $b = \frac{1}{2} \log_4 n = \frac{1}{4} \log_2 n$.
 - Split the input into blocks of size b . Compute an array of minimum values from each block.
 - Build a sparse table on that array of minima.
 - Build per-block RMQ structures for each block, using Cartesian tree numbers to avoid recomputing RMQ structures unnecessarily.
 - Make queries using the standard hybrid solution approach.
- This is an **$\langle O(n), O(1) \rangle$** solution to RMQ!

Wait a Minute...

- This approach assumes that the Cartesian tree numbers will fit into individual machine words!
- If $b = \frac{1}{2}\log_4 n = \frac{1}{4}\log_2 n$, then each Cartesian tree number will have $\frac{1}{2} \log_2 n$ bits.
- Cartesian tree numbers will fit into a machine word if n fits into a machine word.
- In the ***transdichotomous machine model***, we assume the problem size always fits into a machine word.
 - Reasonable – think about how real computers work.
- In some settings this might be too strong of an assumption, but, practically speaking, this is fine!

Practical Concerns

- This structure is both theoretically optimal and fast in practice.
- In their 2005 paper, Fischer and Heun noted that a simple $\langle O(n), O(\log n) \rangle$ hybrid outperforms the $\langle O(n), O(1) \rangle$ solution in practice.
 - Constant factor in the Fischer-Heun's $O(n)$ preprocessing is a bit higher.
 - Constant factor in the hybrid approach's $O(n)$ and $O(\log n)$ are very low.
- Is that still true today? That would be a great question to explore!

The Method of Four Russians

- The technique employed here is an example of the ***Method of Four Russians***.
 - Break the problem of size n into subproblems of size b , plus some top-level problem of size n / b .
 - This is called a ***macro/micro*** decomposition.
 - Solve all possible subproblems of size b .
 - Here, we only solved the subproblems that actually came up in the original array, but that's just an optimization.
 - Solve the overall problem by combining solutions to the micro and macro problems.
- This is a great way to shave off log factors, which are both theoretical and practical wins!

Why Study RMQ?

- I chose RMQ as our first problem for a few reasons:
 - ***See different approaches to the same problem.***
Different intuitions produced different runtimes.
 - ***Build data structures out of other data structures.***
Many modern data structures use other data structures as building blocks, and it's very evident here.
 - ***See the Method of Four Russians.*** This trick looks like magic the first few times you see it and shows up in lots of places.
 - ***Explore modern data structures.*** This is relatively recent data structure (2005), and I wanted to show you that the field is still very active!
- So what's next?

Next Time

- ***Tries***
 - A canonical string data structure.
- ***Suffix Trees***
 - Exposing the hidden structures of strings.