

Range Minimum Query (RMQ)

Input: $A = (A[1], A[2], \dots, A[n])$ array of size n .

Task: preprocess A , so that we can answer queries

$\text{range-min}(i, j) \rightarrow \min_{i \leq k \leq j} \{A[k]\}$

(preferably return index k)

$\rightarrow k$ s.t. $A[k] < A[t]$ for all $i \leq t \leq j$

e.g.

1	2	3	4	5
12	3	6	8	17

$\text{range-min}(2, 4) \rightarrow \textcircled{2} A[2] = 3$

$\text{range-min}(3, 5) \rightarrow \textcircled{3} A[3] = 6$

• static* case (array fixed, no insert/delete)
* some solutions can be made dynamic

• assume array entries are distinct (ow. perturb by random $[0, \epsilon]$)

Quality of an algorithm measured as:

- preprocessing time $p(n)$
- query time $q(n)$

expressed as $\langle p(n), q(n) \rangle$

Best we can hope for: $\langle O(n), O(1) \rangle$

must at least read the array — best possible

(We will get there.)

Observation

If we wanted range-sum, then easy.

$\text{range-sum}(i, j) \rightarrow \text{return } A[i] + \dots + A[j]$

Idea: precompute all prefix sums S

$$S[i] = A[1] + A[2] + \dots + A[i]$$

$$S[0] = 0$$


$$S[1] = A[1]$$

$$S[2] = A[1] + A[2] = S[1] + A[2]$$

$$\dots$$
$$S[n] = S[n-1] + A[n]$$

} Overall $O(n)$ time

range-sum (i, j)
 $\text{return } (S[j] - S[i-1])$
 inverse of + $\langle O(n), O(1) \rangle$
 $O(1)$



range-min more difficult
 (min no inverse)

Why solve RMQ (Applications)



- string algorithms (pattern matching)
- algorithms on trees/graphs
- succinct data structures

Approach 1. Precompute everything

Compute $M[i, j]$ = answer to range-min (i, j)


Naively $O(n^3)$

Dynamic Programming $O(n^2)$



$$M[i, j] = A[i] \quad i \leq j$$

$$M[i, j] = \min \{ M[i, j-1], A[j] \}$$



running time:

$\langle O(n^2), O(1) \rangle$

Approach 1.

range-min (i, j)
 $\text{return } M[i, j]$

Approach 2. Precompute nothing

range-min (i, j)

$m = A[i]$
 for $k = i+1, j$
 $m = \min \{ m, A[k] \}$
 return m

running time:

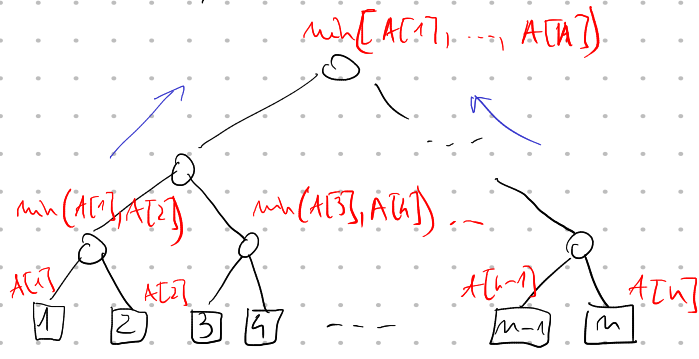
$\langle O(1), O(n) \rangle$

Approach 2.

Approach 3: Augmented trees




Build a balanced BST, using the indices as keys.



- assume $n = 2^k$
- augment nodes with min array entry away from index in subtree (also store index of min)

- tree size $2n-1$
- augmentation can be computed in $O(n)$ time (bottom-to-top level by level)

- answer query range-min (i, j)
 - search for $[i, j]$




- take min over \rightarrow node-minime for all nodes on path with trees happy inside
- \rightarrow and also $[i, j]$ itself
- $O(\log n)$ (depth $\leq \log_2 n + 1$)

overall running time $\langle O(n), O(\log n) \rangle$ Approach 3-

(this is not yet optimal, we want $\langle O(n), O(1) \rangle$)

Detour: an application of RMQ

Find lowest common ancestor (lca) in a tree in $O(1)$ time



- pre-process tree ($O(n)$) s.t. $lca(i, j)$ take $O(1)$ time
- $\langle O(n), O(1) \rangle$ lca
- $lca(h, i) \rightarrow c$
- $lca(g, e) \rightarrow a$

→ useful for distance-queries in trees
(store level/depth of nodes)

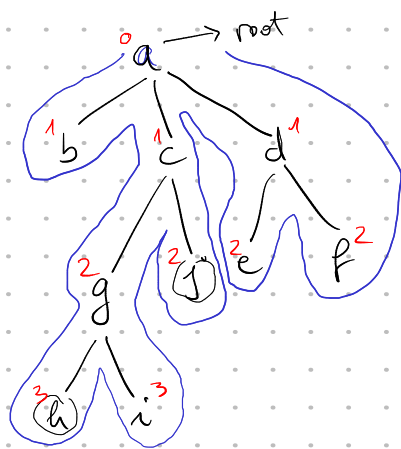
e.g. XML-documents

$$\text{dist}(x, y) = \text{depth}(x) + \text{depth}(y) - 2 \text{depth}(\text{lca}(x, y))$$



Solve LCA problem using RMQ

- Do a "DFS" on the tree, record it as an array.
- RMQ on this array



- for every traversed edge, record both endpoints (Euler-tour)

ET

abacghgigcjcadedfda

→ $Size = 2n - 1$

depths

0 1 0 1 2 3 2 3 2 1 2 1 0 1 2 1 2 1 0

$\text{lca}(x, y) \rightarrow$ return $\text{range-min}(i_x, i_y)$

e.g. $\text{lca}(h, j) \rightarrow c$

$\text{lca}(g, e) \rightarrow a$

minimum
according to
depths \downarrow

index of some
occurrence of x, y
in ET

$n-1$ edges, each traversed twice
each edge-traversal adds one node
+ beginning at root $\Rightarrow 2n-1$ entries

$\langle O(n), O(1) \rangle$ for RMQ \Rightarrow $\langle O(n), O(1) \rangle$ for LCA

(we just showed)

← ? (next)

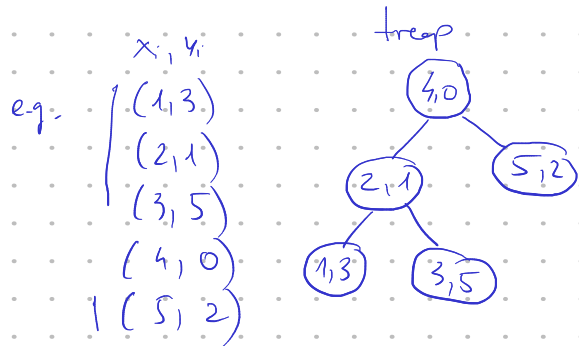
Solve RMQ using LCA

Recall: treap or Cartesian tree

tree with nodes (x_i, y_i) $i = \overline{1, n-1, n}$

s.t. BST according to x_i values

min-heap according to y_i values




Given array $A[1], \dots, A[n]$

"Cartesian tree of array"

Build a heap on pairs $(i, A[i])$
(Cartesian tree)

BST

min-heap



$A[i]$	20	5	8	13	11	2	9	14	6
i	1	2	3	4	5	6	7	8	9

↑ ↑ ↑ ↑ ↑ ↑ ↑ ↑

It follows that

$\langle O(n), O(1) \rangle$ for lca $\Rightarrow \langle O(n), O(1) \rangle$ for RMQ

- build treap
- answer LCA queries on it

Proof (Claim)

- $lca(i, j)$ is between i, j
- subtree of $lca(i, j)$ contains entire interval $[i, j]$
- $lca(i, j)$ has minimum value in all its subtree




minimum value

$A[i], \dots, A[j]$

□

$LCA \Leftrightarrow RMQ$

Approach 4. Divide & Conquer (Blocks)



- split array into $\frac{n}{b}$ blocks of size b
- assume n divisible by b
- precompute minima in blocks
 k^{th} block $\rightarrow \min_k$


range-min (i, j)

- figure out in which block i, j fall

- return minimum of:



min in block of i from i to the end of the block



min in block of j from beginning to j

- (3) • minimum over minima of full blocks between i and j .

- preprocessing time: $O(n)$ \rightarrow computing block minima

- query time: $O(b) + O(\frac{n}{b}) = O(b + \frac{n}{b})$

(if i, j fall in the same block, then easier, then we took \min that block only.)



choose $b = \sqrt{n}$

Query time $O(\sqrt{n})$

Overall: $\langle O(n), O(\sqrt{n}) \rangle$ Approach 4.

Approach 5. Canonical ranges.

Recall $\langle O(n^2), O(1) \rangle$ "precompute all"

Let's try to precompute fewer queries, so that we can still reconstruct range-min(i, j) in $O(1)$ time.

Idea: compute all range-queries of length 2^k

precompute range-min($i, i + 2^k - 1$)

for all $i = \overline{1, n}$

for all $k = \overline{0, \log n}$

\rightarrow Canonical ranges


for i, k such that $i + 2^k - 1 \leq n$

$O(n \log n)$ ranges.

Claim preprocessing takes $O(n \log n)$ time.

Use D-P

- first compute all ranges of length 1
- next compute



2 → each takes $O(1)$

4
8
 2^k

(Computing a range-min for a range of size 2^k can be done as min of two (already computed) of size 2^{k-1})

Queries


range-min (i, j)

- if $j = i + 2^k - 1$, for some k . (length of range is power of 2 - min is precomputed)

return precomputed value.

- else $j - i$ is between $2^k - 1$ and $2^{k+1} - 1$ (for some k)

return $\min \{ \text{range-min}(i, i + 2^k - 1), \text{range-min}(j - 2^k + 1, j) \}$



$O(1)$ time

Interval of interest is covered by two intervals of size 2^k . (minimum is correct, even if overlap)

e.g.




length 2^4 , no precomputed

range-min $(3, 20) \rightarrow \min (\text{range-min}(3, 18), \text{range-min}(5, 20))$

Approach 5.

Overall: $\langle O(n \log n), O(1) \rangle$

Approach 6. Divide & Conquer (Recursively)



$\frac{n}{b}$ blocks of size b

range-min(i, j)

need to compute range-min

- (1) within blocks \rightarrow MICRO
- (2) among block-minima \rightarrow MACRO

let's use recursively same RMA solution for both tasks.

High level scheme

- Split A into $\frac{n}{b}$ blocks of size b
- Compute min in each block $\rightarrow O(n)$
- Build RMA solution within each block with time $\langle P_1, Q_1 \rangle$ — MICRO
- Build RMA solution between block-minima with time $\langle P_2, Q_2 \rangle$ — MACRO

Overall preprocessing time:

$$P(n) = O(n) + P_1(b) \cdot \frac{n}{b} + P_2\left(\frac{n}{b}\right)$$

$$Q(n) = Q_1(b) + Q_2\left(\frac{n}{b}\right)$$

Some concrete constructions

6.a) $b = \sqrt{n}$

$$\left. \begin{aligned} \langle P_1, Q_1 \rangle &= \langle O(1), O(n) \rangle \\ \langle P_2, Q_2 \rangle &= \langle O(1), O(n) \rangle \end{aligned} \right\} \text{"no preproc." Approach 2}$$

$$P(n) = O(n) + O(1) \cdot \frac{n}{b} + O(1) = O(n)$$

$$Q(n) = O(\sqrt{n}) + O(\sqrt{n}) = O(\sqrt{n})$$

$$\langle O(n), O(\sqrt{n}) \rangle \quad \text{6.a)}$$

just like Approach 4.

6. b) $b = \log n$

$\langle p_1, q_1 \rangle = \langle O(1), O(n) \rangle$

"no prepr."

$\langle p_2, q_2 \rangle = \langle O(n \log n), O(1) \rangle$

"Conical rays" appr. 5

$P(n) = O(n) + \frac{n}{\log n} + \frac{n}{\log n} \cdot \log\left(\frac{n}{\log n}\right) \in O(n)$

$q(n) = O(\log n) + O(1) = O(\log n)$

$\langle O(n), O(\log n) \rangle$ 6. b)

(save running time as approach 3)

6. c) $b = \log n$

$\langle p_1, q_1 \rangle = \langle O(n \log n), O(1) \rangle$

$\langle p_2, q_2 \rangle = \langle O(n \log n), O(1) \rangle$

$P(n) = O(n) + \frac{n}{\log n} \cdot \log n \cdot \log \log n + \frac{n}{\log n} \cdot \log\left(\frac{n}{\log n}\right) = O(n \cdot \log \log n)$

$q(n) = O(1)$

$\langle O(n \log \log n), O(1) \rangle$ 6. c)

best so far!

6. d)

use conical rays (appr. 5) for macro

use augmented tree (appr. 3) for micro

$\langle p_1, q_1 \rangle = \langle O(n), O(\log n) \rangle$

$\langle p_2, q_2 \rangle = \langle O(n \log n), O(1) \rangle$

$b = \log n$

$P(n) = O(n) + \frac{n}{b} + \frac{n}{\log n} \cdot \log\left(\frac{n}{\log n}\right) = O(n)$

$q(n) = O(\log \log n) + O(1) = O(\log \log n)$

Bottomline: $\langle O(n), O(\log \log n) \rangle$ 6. d)

also "best so far"

6. e) Phy n recursively


→ Further improvement

need a more idea.

will not get to $\langle O(n), O(1) \rangle$

Approach 7. "All-in"

- Block-decomposition, just one level




$\frac{n}{b}$ blocks of size b

Macro-structure $\langle O(n \log n), O(1) \rangle$ (Canonical ranges approach 5.)

Micro-structure $\rightarrow \langle O(n^2), O(1) \rangle$ (precompute all, but take advantage of blocks that are the same)


We need to build RMA structure for $\frac{n}{b}$ blocks of size b

- some blocks may be the same, then we can reuse same RMA structure



Obs. If two blocks have the same "ordering of elements", then use same RMA micro-structure

Obs. two blocks are "structurally the same" (for every range-min query, same answer)



Caterian trees of two blocks are the same

Q. How many different Caterian trees of size b ?

A. $\# < 4^b$

of different binary trees of size n

is $C_n \rightarrow$ Catalan number.

n	1	2	3	4	5
C_n	1	2	5	14	42

See OEIS.org/A000108 or Wikipedia on "Catalan numbers"

Recurrence:


$$C_n = \sum_{i=1}^n C_{i-1} \cdot C_{n-i}$$

base case:

$$C_0 = 1$$

$$C_n = \frac{1}{n+1} \cdot \binom{2n}{n} < 4^n$$

recall LCA, RMA equivalence. If Caterian trees the same, LCA gives same \Rightarrow RMA same



bijection

binary trees of size n

valid systems of parentheses of length $2n$

$((()())())$

length $2n$

$$\# < 2^{2n} = 4^n$$

(# strings of length $2n$ that encode a valid system of parentheses)

of "structurally different blocks" $< 4^b$

Back to approach 7.


Preprocessing

- choose block size $b = \log_4 \sqrt{n}$

- on each block build Cartesian tree $O(b)$ time, encoding on $\leq 2b$ bits

Save type of block for $i=1, \dots, \frac{n}{b}$

different types $< 4^b = \sqrt{n}$




$1 \leq \text{type}_i \leq \sqrt{n}$

(time for this: $O(n)$)

- For each type of block do $\langle O(n^2), O(n) \rangle$ scheme (full preprocessing)

(total time for this:

$$\sqrt{n} \cdot (\log_4 \sqrt{n})^2 = O(n)$$



RMA structure for each type of block


- within block-minima we use $\langle n \log n, O(n) \rangle$ scheme, we have to preprocess it

time for this:

$$\frac{n}{b} \cdot \log \frac{n}{b} = \frac{n}{\log_4 \sqrt{n}} \cdot \log \frac{n}{\log_4 \sqrt{n}} \leq \frac{n}{\log_4 \sqrt{n}} \cdot \log n = O(n)$$

Total time for all preprocessing: $O(n)$

- Query-time



$O(1)$ time

Overall query time:

$O(n)$

(type of each block saved in preprocessing)

Finally:

$$\langle O(n), O(n) \rangle$$

Approach 7.