Persistent Segment Tree

Oliver-Matis Lill

January 15, 2017

Problem

- We have some initial array with $N \leq 10^5$ elements
- At each of the $T \leq 10^5$ timepoints $1, \ldots, T$, some element i gets updated to a new value x
- We have $Q \leq 10^5$ queries each of which asks for the sum of array elements at some segment $[l, r]$ in some past timepoint t
- Queries are online and asked after the updates

Persistent Segment Tree

Idea

- The queries would be easy to answer if we had a snapshot of the segment tree at each timepoint
- When updating some element, at most $O(\log n)$ nodes in the segment tree get changed: the nodes along the path from root to the updated leaf
- For each timepoint, instead of creating a copy of the entire segment tree, copy only nodes on the path to be updated and update them
- The nodes on the copied path can have children from previous timepoints

Example

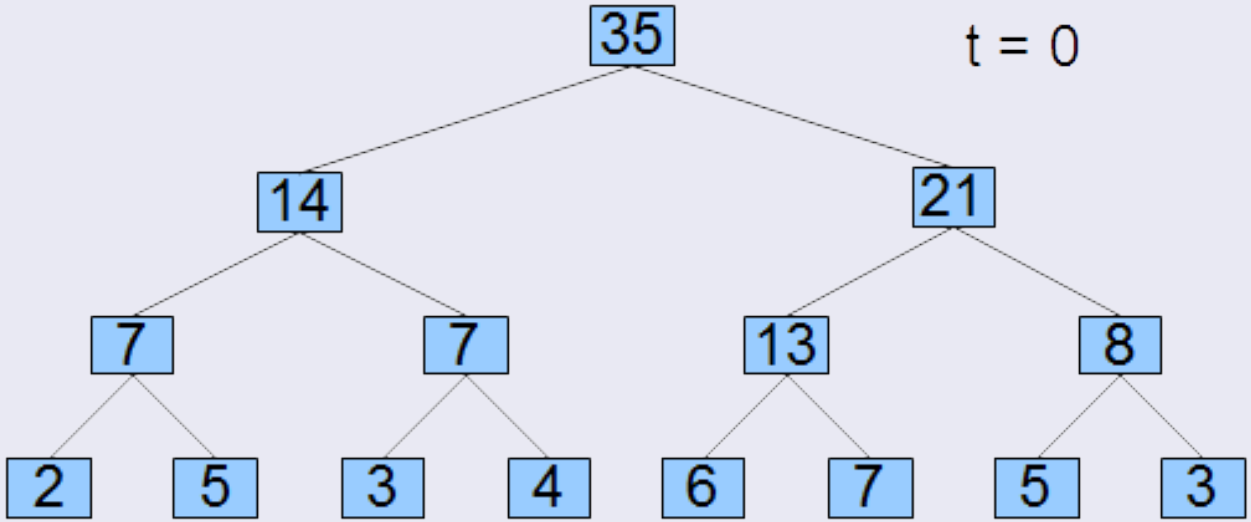

Example

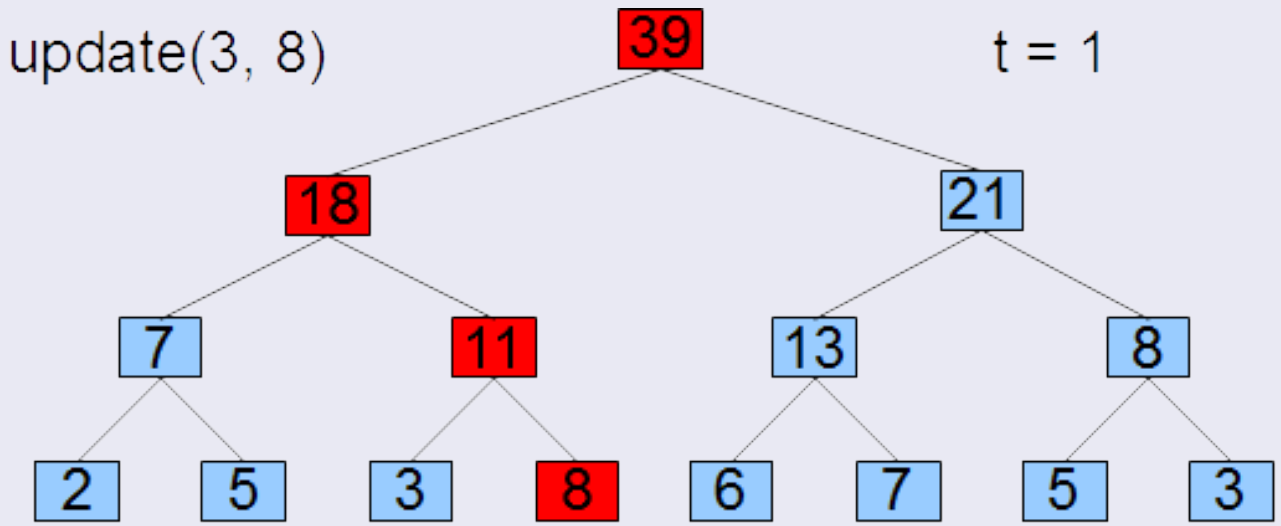

Example

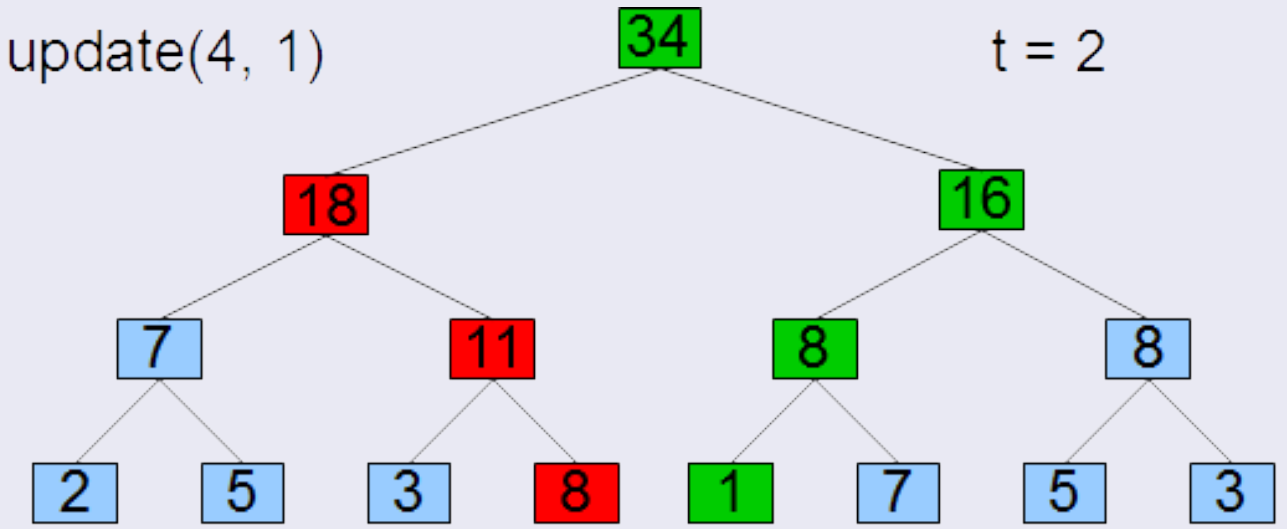

Example

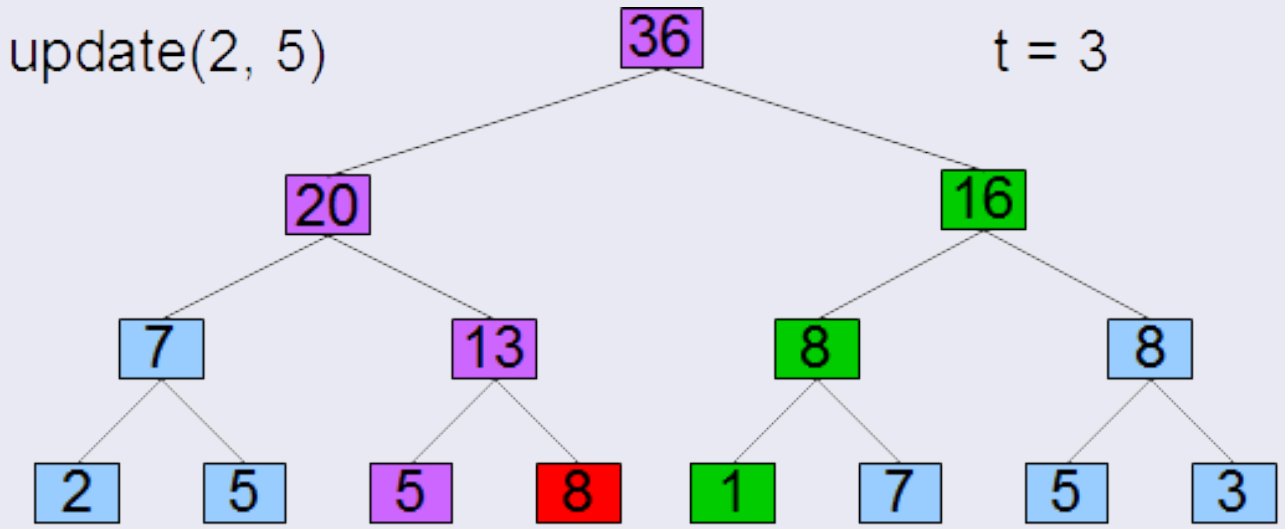

Example

The data stored for each timepoint looks like the following:

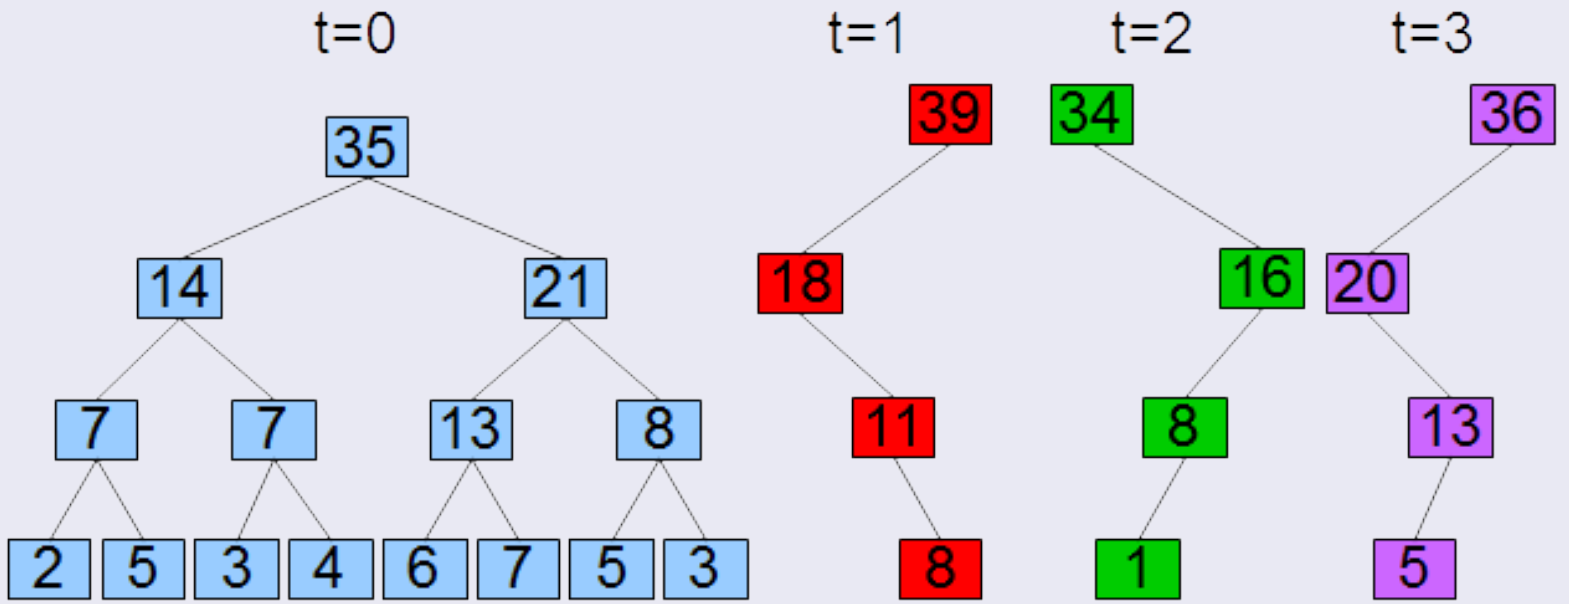

Example

And the entire persistent tree with all parent-child relations is the following:

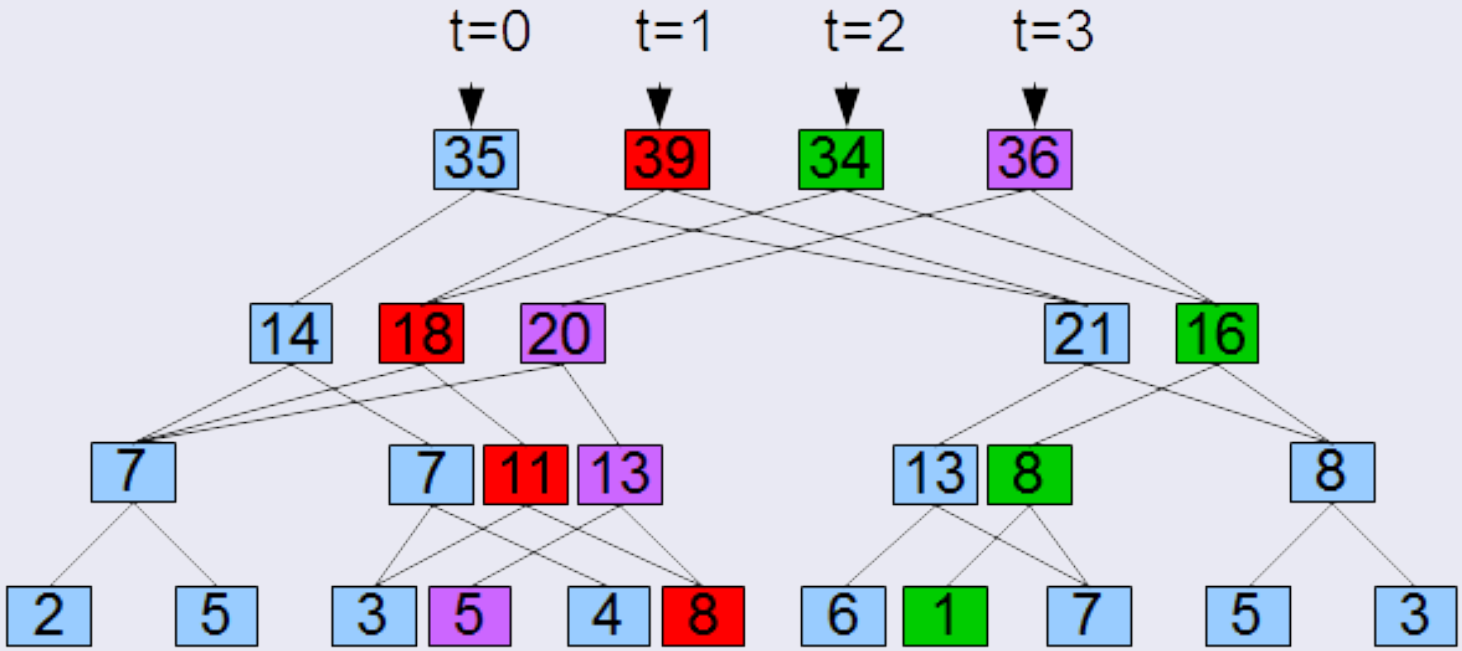

Persistent Segment Tree

- The total storage complexity is $O(N + T \log N)$
- The time complexity for answering each query is $O(\log N)$
- Some problems:
  1. `https://www.codechef.com/problems/SEGSUMQ`
  2. `https://www.hackerrank.com/contests/w22/challenges/sequential-prefix-function`