# Lecture 4: Divide and Conquer: van Emde Boas Trees 

- Series of Improved Data Structures
- Insert, Successor
- Delete
- Space

This lecture is based on personal communication with Michael Bender, 2001.

## Goal

We want to maintain $n$ elements in the range $\{0,1,2, \ldots, u-1\}$ and perform Insert, Delete and Successor operations in $\mathcal{O}(\log \log u)$ time.

- If $n=n^{\mathcal{O}(1)}$ or $n^{(\log n)^{\mathcal{O}(1)}}$, then we have $\mathcal{O}(\log \log n)$ time operations
- Exponentially faster than Balanced Binary Search Trees
- Cooler queries than hashing
- Application: Network Routing Tables
$-u=$ Range of IP Addresses $\rightarrow$ port to send
$\left(u=2^{32}\right.$ in $\left.\operatorname{IPv} 4\right)$

Where might the $\mathcal{O}(\log \log u)$ bound arise ?

- Binary search over $\mathcal{O}(\log u)$ elements
- Recurrences

$$
\begin{aligned}
& -T(\log u)=T\left(\frac{\log u}{2}\right)+\mathcal{O}(1) \\
& -T(u)=T(\sqrt{u})+\mathcal{O}(1)
\end{aligned}
$$

## Improvements

We will develop the van Emde Boas data structure by a series of improvements on a very simple data structure.

## Bit Vector

We maintain a vector $V$ of size $u$ such that $V[x]=1$ if and only if $x$ is in the set. Now, inserts and deletes can be performed by just flipping the corresponding bit in the vector. However, successor/predecessor requires us to traverse through the vector to find the next 1-bit.

- Insert/Delete: $\mathcal{O}(1)$
- Successor/Predecessor: $\mathcal{O}(u)$

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :---: | :---: |
| 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 1 | 0 | 0 | 0 | 0 | 1 |

Figure 1: Bit vector for $u=16$. THe current set is $\{1,9,10,15\}$.

## Split Universe into Clusters

We can improve performance by splitting up the range $\{0,1,2, \ldots, u-1\}$ into $\sqrt{u}$ clusters of size $\sqrt{u}$. If $x=i \sqrt{u}+j$, then $V[x]=V$.Cluster $[i][j]$.

$$
\begin{aligned}
\operatorname{low}(x) & =x \bmod \sqrt{u}=j \\
\operatorname{high}(x) & =\left\lfloor\frac{x}{\sqrt{u}}\right\rfloor=i \\
\text { index }(i, j) & =i \sqrt{u}+j
\end{aligned}
$$



Figure 2: Bit vector $(u=16)$ split into $\sqrt{16}=4$ clusters of size 4.

- Insert:
- Set V.cluster $[\operatorname{high}(x)][\operatorname{low}(x)]=1$
- Mark cluster high $(x)$ as non-empty
- Successor:
- Look within cluster high( $x$ )

$$
\begin{array}{r}
\mathcal{O}(\sqrt{u}) \\
\mathcal{O}(\sqrt{u}) \\
\text { Total }=\frac{\mathcal{O}(\sqrt{u})}{\mathcal{O}(\sqrt{u})}
\end{array}
$$

- Find minimum entry $j$ in that cluster
- Return index $(i, j)$


## Recurse

The three operations in Successor are also Successor calls to vectors of size $\sqrt{u}$. We can use recursion to speed things up.

- V.cluster $[i]$ is a size- $\sqrt{u}$ van Emde Boas structure $(\forall 0 \leq i<\sqrt{u})$
- V.summary is a size- $\sqrt{u}$ van Emde Boas structure
- V.summary $[i]$ indicates whether $V . c l u s t e r[i]$ is nonempty
$\operatorname{INSERT}(V, x)$
$1 \operatorname{Insert}(V . \operatorname{cluster}[\operatorname{high}(x)]$, low $[x])$
$2 \operatorname{Insert}(V . s u m m a r y$, high $[x])$
So, we get the recurrence:

$$
\begin{aligned}
T(u) & =2 T(\sqrt{u})+\mathcal{O}(1) \\
T^{\prime}(\log u) & =2 T^{\prime}\left(\frac{\log u}{2}\right)+\mathcal{O}(1) \\
\Longrightarrow T(u) & =T^{\prime}(\log u)=\mathcal{O}(\log u)
\end{aligned}
$$

$\operatorname{SUCCESSOR}(V, x)$
$1 i=\operatorname{high}(x)$
$2 j=\operatorname{Successor}(V . c l u s t e r[i], j)$
3 if $j==\infty$
$4 \quad i=\operatorname{Successor}$ (V.summary,$i$ )
$5 \quad j=\operatorname{Successor}($ V.cluster $[i],-\infty)$
6 return index $(i, j)$

$$
\begin{aligned}
T(u) & =3 T(\sqrt{u})+\mathcal{O}(1) \\
T^{\prime}(\log u) & =3 T^{\prime}\left(\frac{\log u}{2}\right)+\mathcal{O}(1) \\
\Longrightarrow T(u) & =T^{\prime}(\log u)=\mathcal{O}\left((\log u)^{\log 3}\right) \approx \mathcal{O}\left((\log u)^{1.585}\right)
\end{aligned}
$$

To obtain the $\mathcal{O}(\log \log u)$ running time, we need to reduce the number of recursions to one.

## Maintain Min and Max

We store the minimum and maximum entry in each structure. This gives an $\mathcal{O}(1)$ time overhead for each Insert operation.
$\operatorname{SUCCESSOR}(V, x)$
$i=\operatorname{high}(x)$
2 if $\operatorname{low}(x)<V$.cluster $[i]$.max
$3 \quad j=\operatorname{Successor}(V . c l u s t e r[i]$, low $(x))$
4 else $i=\operatorname{Successor}(V . s u m m a r y, \operatorname{high}(x))$
$5 \quad j=$ V.cluster $[i]$.min
6 return index $(i, j)$

$$
\begin{aligned}
T(u) & =T(\sqrt{u})+\mathcal{O}(1) \\
\Longrightarrow T(u) & =\mathcal{O}(\log \log u)
\end{aligned}
$$

## Don't store Min recursively

The Successor call now needs to check for the min separately.

$$
\begin{equation*}
\text { if } x<V \text {.min }: \text { return } V \text {.min } \tag{1}
\end{equation*}
$$

$\operatorname{INSERT}(V, x)$

```
if \(V . \min ==\) None
    \(V . \min =V . \max =x \quad \triangleright \mathcal{O}(1)\) time
    return
if \(x<V\).min
    \(\operatorname{swap}(x \leftrightarrow V\).min \()\)
if \(x>V\).max
    \(V \cdot \max =x)\)
if \(V\).cluster \([\operatorname{high}(x)==\) None
    Insert(V.summary,high \((x)) \quad \triangleright\) First Call
\(\operatorname{Insert}(V . \operatorname{cluster}[\operatorname{high}(x)]\),low \((x)) \quad \triangleright\) Second Call
```

If the first call is executed, the second call only takes $\mathcal{O}(1)$ time. So

$$
\begin{aligned}
T(u) & =T(\sqrt{u})+\mathcal{O}(1) \\
\Longrightarrow \quad T(u) & =\mathcal{O}(\log \log u)
\end{aligned}
$$

$\operatorname{Delete}(V, x)$

```
if \(x==V\).min \(\quad \triangleright\) Find new min
    \(i=\) V.summary.min
    if \(i=\) None
        \(V . \min =V . \max =\) None \(\quad \triangleright \mathcal{O}(1)\) time
        return
    \(V\). min \(=\) index \((i, V . c l u s t e r[i] . m i n) \quad \triangleright\) Unstore new min
Delete(V.cluster[high(x)],low(x)) \(\triangleright\) First Call
if V.cluster \([\) high \((x)] . \min ==\) None
    Delete(V.summary,high \((x)) \quad \triangleright\) Second Call
\(\triangleright\) Now we update \(V\).max
    if \(x==V\).max
if \(V\).summary.max \(=\) None
else
    \(i=V\). summary.max
    \(V . \max =\) index \((i, V . \operatorname{cluster}[i] . \max )\)
```

If the second call is executed, the first call only takes $\mathcal{O}(1)$ time. So

$$
\begin{aligned}
T(u) & =T(\sqrt{u})+\mathcal{O}(1) \\
T(u) & =\mathcal{O}(\log \log u)
\end{aligned}
$$

## Lower Bound [Patrascu \& Thorup 2007]

Even for static queries (no Insert/Delete)

- $\Omega(\log \log u)$ time per query for $u=n^{(\log n)^{\mathcal{O}(1)}}$
- $\mathcal{O}(n \cdot \operatorname{poly}(\log n))$ space


## Space Improvements

We can improve from $\Theta(u)$ to $\mathcal{O}(n \log \log u)$.

- Only create nonempty clusters
- If $V$.min becomes None, deallocate $V$
- Store V.cluster as a hashtable of nonempty clusters
- Each insert may create a new structure $\Theta(\log \log u)$ times (each empty insert)
- Can actually happen [Vladimir Čunát]
- Charge pointer to structure (and associated hash table entry) to the structure This gives us $\mathcal{O}(n \log \log u)$ space (but randomized).


## Indirection

We can further reduce to $\mathcal{O}(n)$ space.

- Store vEB structure with $n=\mathcal{O}(\log \log u)$ using BST or even an array
$\Longrightarrow \mathcal{O}(\log \log n)$ time once in base case
- We use $\mathcal{O}(n / \log \log u)$ such structures (disjoint)
$\Longrightarrow \mathcal{O}\left(\frac{n}{\log \log u} \cdot \log \log u\right)=\mathcal{O}(n)$ space for small
- Larger structures "store" pointers to them

$$
\Longrightarrow \mathcal{O}\left(\frac{n}{\log \log u} \cdot \log \log u\right)=\mathcal{O}(n) \text { space for large }
$$

- Details: Split/Merge small structures

MIT OpenCourseWare
http://ocw.mit.edu

### 6.046J / 18.410J Design and Analysis of Algorithms

Spring 2015

For information about citing these materials or our Terms of Use, visit: http://ocw.mit.edu/terms.

