

Lower Bounds for Predecessor Search



MADALGO Summer School 2010
Monday, Afternoon I

Results

[van Emde Boas FOCs'75]
Predecessor search in $O(n)$ space, $O(\lg \lg u)$ time

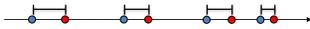
[Pătrașcu, Thorup STOC'06] Let $n^{1+\epsilon} \leq u \leq n^{\text{polylog}(n)}$
Any data structure of size $n \cdot \log^{O(1)} n$ requires time $\Omega(\lg \lg u)$

In fact, n/u /space trade-off completely understood ☺

$$\min \begin{cases} \log_{u/n} n \\ \lg \frac{u - \lg n}{u} \\ \frac{\lg \frac{u}{n}}{\lg(\frac{u}{n} - \lg \frac{u}{n})} \\ \frac{\lg \frac{u}{n}}{\lg(\frac{u}{n} - \frac{u}{u - \lg n})} \end{cases}$$

Colored Predecessor

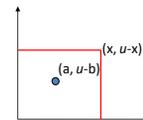
Color the set red-blue $S = \{x_1, x_2, x_3, x_4, x_5, \dots\}$
Lower bound holds even for finding color of predecessor
 \Rightarrow colored predecessor \equiv predecessor
 \Rightarrow existential static 1D stabbing \equiv predecessor



Query stabs a segment \Leftrightarrow predecessor is blue

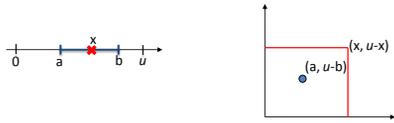
Colored Predecessor

Color the set red-blue $S = \{x_1, x_2, x_3, x_4, x_5, \dots\}$
Lower bound holds even for finding color of predecessor
 \Rightarrow colored predecessor \equiv predecessor
 \Rightarrow existential static 1D stabbing \equiv predecessor
 \Rightarrow static existential 2D dominance \equiv predecessor

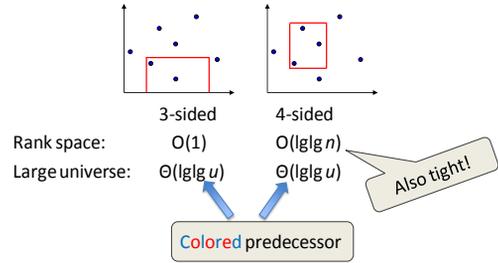


Colored Predecessor

In general: stabbing in d dimensions reduces to dominance queries in $2d$ dimensions

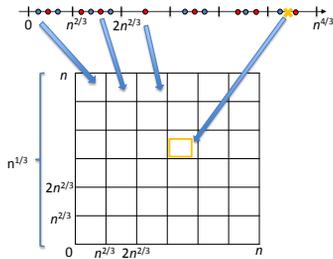


Review: 2D Range Queries



4-sided Queries in Rank Space

Predecessor search for $u=n^{4/3}$ requires time $\Omega(\lg n)$



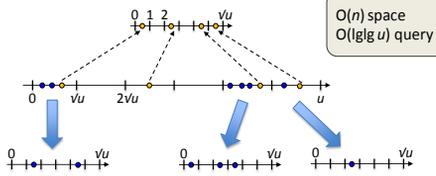
Time for...

The Lower Bound Proof

Let $u = n^3$
 Any data structure of size $O(n)$ requires time $\Omega(\lg n)$

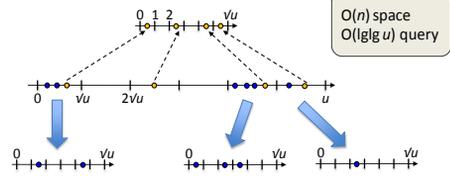
Review: van Emde Boas

$\text{pred}(q, S)$: if $\lfloor q/vu \rfloor \in \text{hash table}$,
 return $\text{pred}(q \bmod vu, \text{bottom structure})$
else return $\text{pred}(\lfloor q/vu \rfloor, \text{top structure})$



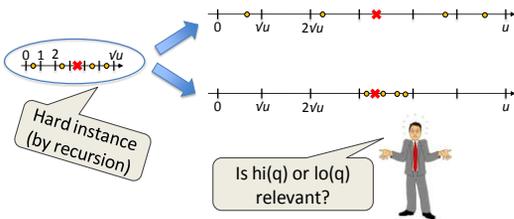
Review: van Emde Boas

$\text{pred}(q, S)$: if $\text{hi}(q) \in \text{hash table}$,
 return $\text{pred}(\text{lo}(q), \text{bottom structure})$
else return $\text{pred}(\text{hi}(q), \text{top structure})$



The Hard Instance

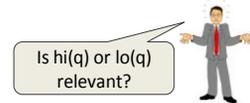
Intuition: In one memory probe,
 one can at most reduce the universe $u \mapsto vu$



The First Memory Access

van Emde Boas always tries top first \Rightarrow read address $f(\text{hi}(q))$
 Can easily be fooled (put all elements in a vu interval)

But I can access $O(n)$ memory locations!
 I'll choose among $O(vn)$ options depending $\text{hi}(q)$
 and $O(vn)$ options depending on $\text{lo}(q)$



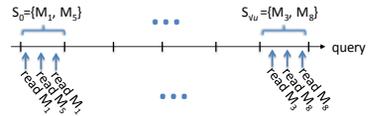
Updating our Intuition

Simple intuition: In one memory probe, one can at most reduce the universe $u \mapsto \sqrt{u}$

Refined intuition:

- either $\text{pred}(\text{hi}(q), \text{hi}(S))$ or $\text{pred}(\text{lo}(q), \text{lo}(S))$ is relevant
- query algorithm can read $\leq \sqrt{vn}$ cells as $f(\text{relevant part})$

Formally

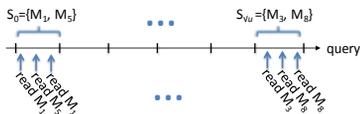


If $(\exists)k \quad |S_k| \leq \sqrt{vn}$:

- place query & data set in segment k
- 1st memory access = $f(\text{lo}(q)) \in S_k$

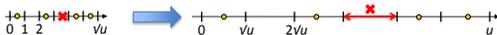


Formally



Otherwise $(\forall)k \quad |S_k| \geq \sqrt{vn}$:

- choose $T = \{O(\sqrt{vn} \cdot \lg n) \text{ cells}\} \Rightarrow$ each S_k is hit
- 1st memory access = $f(\text{hi}(q), \text{lo}(q)) \in S_{\text{lo}(q)}$
- make $\text{lo}(q)$ irrelevant \Rightarrow fix to make $f(\text{hi}(q), *) \in T$



What Did We Prove?

If there exists a solution to $\text{Pred}(n, u)$ with:

- space complexity: $O(n)$
- query complexity: t memory reads

\Rightarrow

There exists a solution to $\text{Pred}(n, \sqrt{u})$ with:

- space complexity: $O(n)$
- $O(\sqrt{vn} \cdot \lg n)$ "public cells"
- query complexity: $t-1$ memory reads

... can be read free of charge

Induction?

Induction:

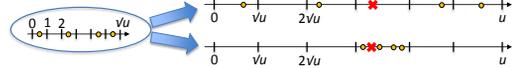
$\text{Pred}(n, u) \mapsto \text{Pred}(n, vu) \mapsto \text{Pred}(n, u^k) \mapsto \dots \mapsto \text{Pred}(n, u^{1/t})$
 time $t \mapsto$ time $t-1 \mapsto$ time $t-2 \mapsto \dots \mapsto$ time 0

Base case: If $u^{1/t} \geq 2$, cannot solve with no memory access! *

* Except $o(n)$ public bits.

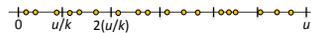
Dealing with Public Bits

Hardness based on hiding **one bit**:



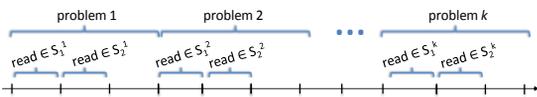
In 2nd round, there are $O(vn \cdot \lg^2 n)$ public bits about the input!

New idea: $\text{Pred}(n, u) = k \times \text{Pred}(n/k, u/k)$



$k \gg vn \cdot \lg^2 n \Rightarrow$ With $O(vn \cdot \lg^2 n)$ public bits, most sub-problems are still hard.

New Induction Plan

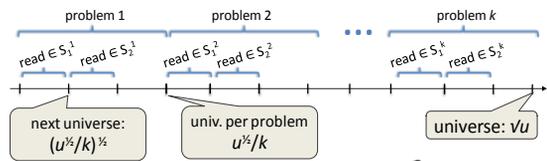


Main Lemma: Fix algorithm to read from a set of $(nk)^{1/2}$ cells
 "Proof":

- problem j is nice if $(\exists)\alpha: |S_\alpha^j| \leq (n/k)^{1/2}$
 \Rightarrow fix hi-part in problem j to α
- problem j is not nice if $(\forall)\alpha: |S_\alpha^j| > (n/k)^{1/2}$
 \Rightarrow choose T to hit all such S_α^j

$k \cdot (n/k)^{1/2} = (nk)^{1/2}$
 $|T| \approx n / (n/k)^{1/2} = (nk)^{1/2}$

New Induction Plan



Set $u=n^3$ and $k = n^k$

	Time	Public bits
$\text{Pred}(n, n^3)$	t	0
$\mapsto \text{Pred}(n, n^{1.5}) \mapsto n^k \times \text{Pred}(n^k, n^k)$	$t-1$	vn
$\mapsto n^k \times \text{Pred}(n^k, n^{3/8}) \mapsto \dots$	$t-2$	$n^{7/8}$

$\Omega(\lg \lg n)$ times ☺

Main Lemma: “Proof” \mapsto Proof

Main Lemma: Fix algorithm to read from a set of $(nk)^{1/2}$ cells

- problem j is nice if $(\exists)\alpha: |S_\alpha^j| \leq (n/k)^{1/2}$
 \Rightarrow fix hi-part in problem j to α
- problem j is not nice if $(\forall)\alpha: |S_\alpha^j| > (n/k)^{1/2}$
 \Rightarrow choose T to hit all such S_α^j

But: Public bits = $f(\text{database})$

1st cell read by query = $f(\text{public bits})$

Look at what query does \Rightarrow Fix some problem adversarially
 \Rightarrow public bits change \Rightarrow Query does something else!

Main Lemma: “Proof” \mapsto Proof

New claim. We can publish $(nk)^{1/2}$ cells such that:

$\Pr[\text{random query reads a published cell}] \geq 1/100$

Induction: If initial query time $< (\lg u)/100$

\Rightarrow at the end $E[\text{query time}] < 0 \Rightarrow$ contradiction

Proof:

- Publish random sample $T = \{ (nk)^{1/2} \text{ cells} \}$
- For each problem j where $\text{lo}(q)$ is relevant (fixed $\text{hi}=\alpha$)
publish S_α^j only if $|S_\alpha^j| \leq (n/k)^{1/2}$

But why does it work?

An Encoding Argument

Impossible: Encode $A[1..k] \in \{0,1\}^k$ with less than k bits on avg.

Proof by contradiction:

- Assume $\Pr[\text{random query reads a published cell}] < 1/100$
- Use this data structure to create an impossible encoder.

The Encoding

1. Generate one random query/subproblem (q_1, q_2, \dots, q_k)
2. Generate random database depending on queries:
 $A[j]=0 \Rightarrow \text{lo}(q_j)$ is relevant in problem j
 $A[j]=1 \Rightarrow \text{hi}(q_j)$ is relevant in problem j
3. Write public bits = $f(\text{database}) \rightarrow o(k)$ bits
4. Classify queries: when $|S_{\text{hi}(q_j)}^j| \leq (n/k)^{1/2}$ query is \odot iff $A[j]=0$
when $|S_{\text{hi}(q_j)}^j| > (n/k)^{1/2}$ query is \ominus iff $A[j]=1$
5. By assumption, $E[\text{number of } \odot \text{ queries}] \geq 99\% k$
So decoder can learn 99% of $A[1..k]$ from public bits!
6. Write in encoding which guesses are wrong
 $\rightarrow \lg(k \text{ choose } k/100) \ll k$ bits