Symbol Table, Dictionary

- records with keys
- INSERT
- SEARCH

Goal: Symbol table implementation

- with O(IgN) GUARANTEED performance
- for both search and insert
- (and other ST operations)


## Three approaches

1. PROBABILISTIC 'guarantee'
2. AMORTIZED "guarantec"
3. WORST-CASE GUARANTEE

## Randomized BSTs

IDEA: new node should be root with probability $I /(N+1)$ DO IT!

```
link insertR(link \(h\), Item item)
        \{ Key \(v=\) key (item), \(t=\) key (h->item);
            if (h == \(z\) ) return NEW (item, \(z, z, 1)\)
        if (rand () < RAND_MAX/( \(h->N+1\) ))
            return insertT(h, item);
        if less ( \(\mathrm{v}, \mathrm{t}\) ) \(\mathrm{h} \rightarrow>1=\) insertR( \(\mathrm{h}->1\), item);
            else \(h \rightarrow>=\) insertR(h->r, item);
        (h->N) ++; return \(h\);
        \}
    void STinsert (Item item)
        \{ head = insertR(head, item); \}
```

Trees have same shape as random BSTS FOR ALL INPUTS Random BSTs: exponentially small chance of bad balance

Insert keys in order: tree shape still random!


## Other operations in randomized BSTS

FIND $k$ th largest

- another use of size field already there JOIN disjoint STs
- straightforward recursive implementation
- to join STs $A$ (of size M) and B (of size N) use $A$ root with probability $M /(M+N)$ use $B$ root with probability $N /(M+N)$ join other tree with subtree recursively
DELETE
- remove the node, do join (above)

THM: Trees still random after delete (!

## Randomized BSTs

Always look like random BSTs


- implementation straightforward
- support all symbol-table ADT opS
- O(log $N$ ) average case
- bad cases provably unlikely


## Skip lists

Idea: Add links to linked-list nodes to make "fast tracks" (


Challenges (see section 13.5 for details):

- how to maintain structure under insertion
- how many links in a particular node?

Bottom line: similar to randomized BSTs

- plus: casier to understand
- minus: more pointer-chasing

Idea: slight modification to root insertion Check two links above current node Orientations differ: same as root insertion Orientations match: do top rotation first













## Splay tree balance

THM: Splay rotations halve the search path

guaranteed performance over SEQUENCE of operations

```
link splay(link h, Item item)
{ Key v = key(item);
    if (h == z) return NEW(item, z, z, 1);
    if (less(v, key(h->item)))
    {
        if (hl == z) return NEW(item, z,h,h->N+1);
        if (less(v, key(hl->item)))
            { hll = splay(hll, item); h = rotR(h); }
        else
            { hlr = splay(hlr, item); hl = rotL(hl);}
        return rotR(h);
    }
    else
    {
        if (hr == z) return NEW(item, h, z, h->N+1)
        if (less(key(hr->item), v))
            { hrr = splay(hrr, item); h = rotL(h); }
        else
            { hrl = splay(hrl, item); hr = rotR(hr);}
        return rotL(h);
    }
}
```


## 2-3-4 trecs

## Top-down 2-3-4 trec construction

Allow one, two, or three keys per node Keep link for every interval beteen keys - 2-node: one key, two children

- 3-node: two keys, three children
- 4-node: three keys, four children



## SEARCH

- compare search key against keys in node
- find interval containing search key
- follow associated link (recursively) INSERT
- search to bottom for key
- 2-node at bottom: convert to a 3-node
- 3-node at bottom: convert to a 4-node
-4-node at bottom: ??




## In top-down 2-3-4 trees,

- all paths from top to bottom are the same length



## Tree height:

- worst case: lgN (all 2-nodes)
-best case: $\operatorname{lgN} / 2$ (all 4-nodes)
- between 10 and 20 for a million nodes
- between 15 and $3^{\circ}$ for a billion nodes

Comparisons within nodes not accounted for

Top-down 2-3-4 tree implementation

Fantasy code (sketch):

```
link insertR(link h, Item item)
    { Key v = key(item);
    link x = h;
    while (x != z)
            { x = therightlink(x, v);
                if fourNode(x) then split(x); }
    if twoNode(x) then makeThree(x, v); else
    if threeNode(x) then makeFour(x, v); else
        return head;
    }
```

Direct implementation complicated because of

- "therightlink( $x, v$ )
- maintaining multiple node types
- large number of cases for "split"
search also more complicated than for BST


## Represent 2-3-4 trees as binary trees

- with 'internal" edges for $3^{-}$and $4^{-n o d e s}$


Correspondence between 2-3-4 and RB trees


Not 1-1 because 3-nodes swing either way

## Splitting nodes in red-black trees

Two cases are casy (need only to switch colors)






Two cases require ROTATIONS



## Red-black tree implementation







A)


```
link RBinsert(link h, Item item, int sw)
```

link RBinsert(link h, Item item, int sw)
{ Key v = key(item);
{ Key v = key(item);
if (h == z) return NEW(item, z, z, 1, 1);
if (h == z) return NEW(item, z, z, 1, 1);
if ((hl->red) \&\& (hr->red))
if ((hl->red) \&\& (hr->red))
{ h->>red = 1; hl->red = 0; hr->>red = 0; }
{ h->>red = 1; hl->red = 0; hr->>red = 0; }
if (less(v, key(h->item)))
if (less(v, key(h->item)))
{
{
hl = RBinsert (hl, item, 0);
hl = RBinsert (hl, item, 0);
if (h->red \&\& hl->red \&\& sw) h = rotR(h);
if (h->red \&\& hl->red \&\& sw) h = rotR(h);
if (hl->red \&\& hll->red)
if (hl->red \&\& hll->red)
{h= rotR(h); h->>red = 0; hr>>red = 1; }
{h= rotR(h); h->>red = 0; hr>>red = 1; }
}
}
else
else
hr = RBinsert(hr, item, 1);
hr = RBinsert(hr, item, 1);
if (h->red \&\& hr->red \&\& !sw) h = rotL(h)
if (h->red \&\& hr->red \&\& !sw) h = rotL(h)
if (hr->red \&\& hrr->red)
if (hr->red \&\& hrr->red)
{h= rotL(h); h->>red = 0; hl>>red = 1; }
{h= rotL(h); h->>red = 0; hl>>red = 1; }
}
}
return h;
return h;
}
}
void STinsert(Item item)
void STinsert(Item item)
{ head=RBinsert (head,item,0); head->red=0; }
{ head=RBinsert (head,item,0); head->red=0; }
{

```
        {
```


## Balance in red-black trees

In red-black trees,

- LONGEST path at most twice as long as SHORTEST path

worst case: less than $2 \operatorname{lgN}$

Comparisons within nodes *are* counted

## Generalize 2-3-4 trees: up to $M$ links per node

Split full nodes on the way down

## Red-black abstraction still works

- BUT might use binary search instead of internal links


## B-trees for external search

- node size $=$ page size
- typical: $M=1000, N<1,000,000,000,000$

Main advantage: flexibility to do fast insert/delete

## Space-time tradeoff

- M large: only a few levels in tree
- M small: less wasted space

Bottom line:

- log_M N page accesses (3 or 4 in practice)


B tree growth


GOAL: ST implementation with $O(\operatorname{lgN})$ GUARANTEE for all ops probabilistic guarantec: random BSTs, skip lists
amortized guarantee: splay trees
optimal guarantec: red-black trees
Algorithms are varations on a theme (rotations when inserting)
Different abstractions, but equivalent
Ex: skip-list representation of 2-3-4 tree


Are balanced trees OPTIMAL?

- worst-case: no (can get ClgN for C>)
- average-case: open

Abstraction extends to give search algs for huge files

- B-trees

