

6.851

Lecture 1

Feb. 7, 2012

Prof. Erik Demaine

TAs: Tom Morgan & Justin Zhang

## TOPICS:


- time travel: remembering/changing the past [THIS WEEK]
- geometry:  $>1$  dimension (maps, DB tables)
- dynamic optimality: is there one best BST?
- memory hierarchy: minimize cache misses
- hashing: most used DS in CS
- integers: beat  $\lg n$  time/op, or prove impossible
- dynamic graphs: changing computer/social network
- strings: search for phrase in text (DNA, web)
- succinct: reduce space to  $\approx$  bare minimum

## Administration:

- video recording of lectures
- requirements: attending lecture,  $\approx$  weekly psets, scribing, project
- signup sheet
- listeners welcome
- problem session (starting  $\sim$  week 3)
- [- scribe for today]

# Theme in this class: THE MODEL MATTERS

## Pointer machine: model of computation



- field = data item OR pointer to node

- operations:  $O(1)$  time each

-  $x = \text{new node}$

-  $x = y.\text{field}$

-  $x.\text{field} = y$

-  $x = y + z$  etc. (data operations)

[ - destroy  $x$  (if no pointers to it) ]

where  $x, y, z$  are fields of root (or root)

$\Rightarrow$  constant working space

e.g. linked list, binary search tree (BST),  
most object-oriented programs

# Temporal data structures:

- persistence [L1]
- retroactivity [L2]

think:  
time travel

## Persistence:

- keep all versions of DS
- DS operations relative to specified version
- update creates (& returns) new version  
(never modify a version)

most of Terminator/  
Sarah Conner Chron.

branching-  
universe model

- 4 levels:

### ① partial persistence:



- update only latest version  
⇒ versions linearly ordered

### ② full persistence:



- update any version  
⇒ versions form a tree

### ③ confluent persistence:



- can combine >1 given version into new v.  
⇒ versions form a DAG

### ④ functional:

- never modify nodes; only create new
- version of DS represented by pointer

movie  
Déjà Vu  
part 1  
Déjà Vu  
part 2

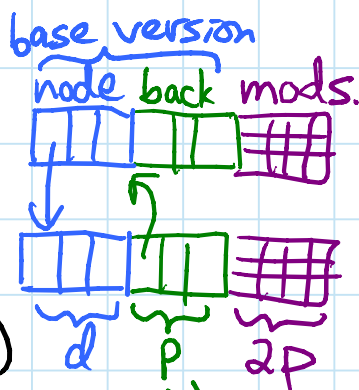
Pullman's book  
Subtle Knife?

TV show Sliders  
movie Primer?

Partial persistence: [Driscoll, Sarnak, Sleator, Tarjan - JCSS 1989]  
 any pointer-machine DS with  $\leq p = O(1)$  pointers to any node (in any version)  
 can be made partially persistent  
 with  $O(1)$  amortized multiplicative overhead  
 &  $O(1)$  space per change

Proof:

- store reverse pointers for nodes in latest version (only)
- allow  $\leq 2p$  (version, field, value) mods. in a node (using that  $p = O(1)$ )
- to read node.field at version  $v$ , check for mods with time  $\leq v$
- when update changes node.field =  $x$ :
  - if node not full: add mod. (now, field,  $x$ )
  - else: - create node' = node with mods. applied



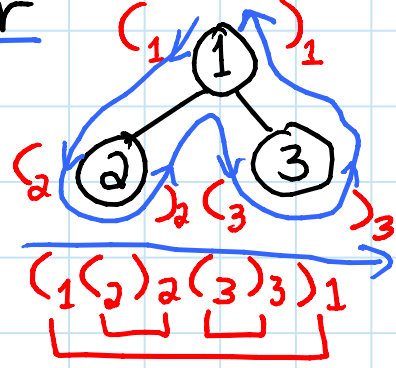
root node  
 part of  
 returned  
 version

- change back pointers to node  $\rightarrow$  node'  
 ↳ found by following pointers
- recursively change pointers to node  $\rightarrow$  ' found via back pointers

(- add back pointer from  $x$  to node)  
 - potential  $\Phi = c \cdot \sum \# \text{mods. in nodes in latest version}$   
 $\Rightarrow$  amortized cost  $\leq \underbrace{c}_{\text{compute mod.}} + \underbrace{c - 2cp}_{\text{if recurse}} + p \text{ recursions}$   
 $\leq 2c. \quad \square$

# Full persistence: ditto [Driscoll et al. 1989]

- linearize tree of versions via Euler tour, marking begin & end of each subtree



- paren sequence = linearized times

- time  $(i$  makes change  $i$

- time  $)i$  unmakes change  $i$

$\Rightarrow$  version  $i$  accumulates changes at times  $\leq (i$

- store times in order-maintenance DS:

[L8: Dietz & Sleator - STOC 1987]

- insert item before/after specified item (like linked list)

- relative order of 2 items?  $< \text{ or } >$

in  $O(1)$  time/op.

- update to version  $i$  represented by 2 mods at times  $(i$  &  $)i$  inserted after  $(i$

update  $\leftarrow$

$\rightarrow$  undo update or before  $)i$



- allow  $\leq 2(d+p+1)$  mods. per node

- forward & backward pointers now symmetric

- when node overflows  $\rightarrow 2(d+p+1)+1$  mods. split into 2 nodes, each  $\approx$  half full (like B-tree)

- find median time  $t^*$

in sorted order of mod. times

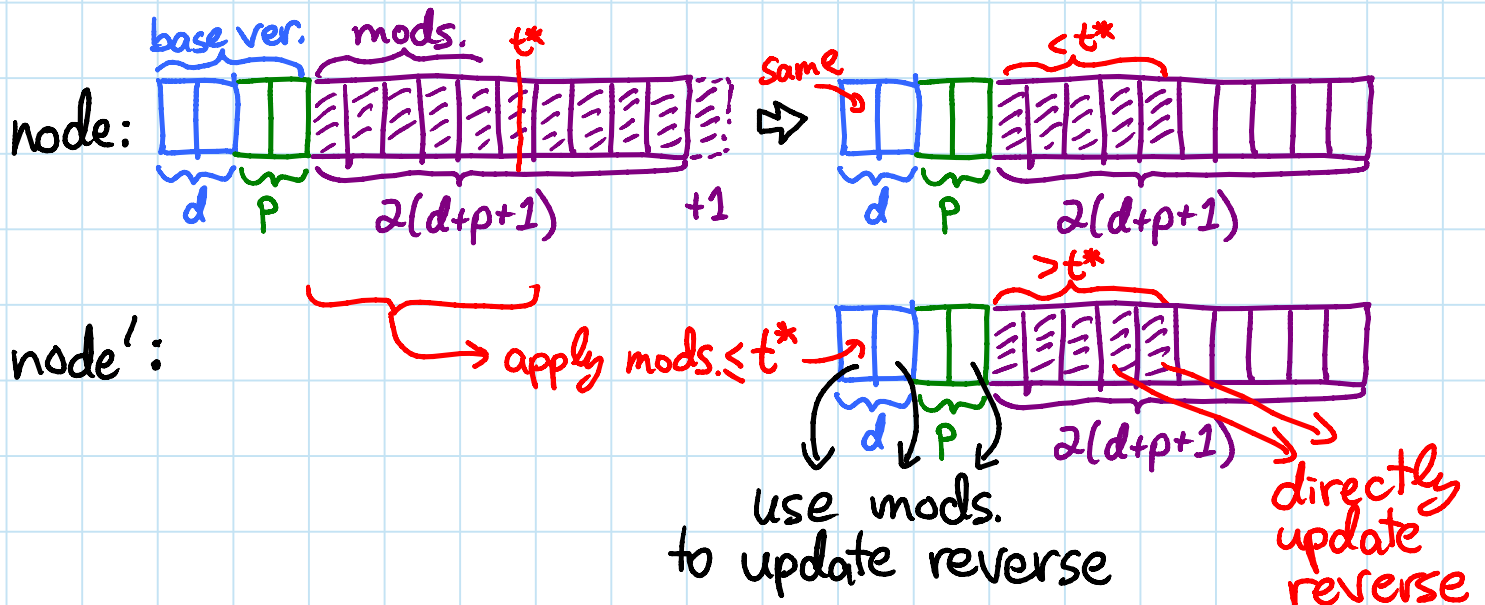


multiple mods. at same time  $t^*$

- old node keeps mods. at  $\leq t^*$  } exclude  $t^*$   
- new node keeps mods. at  $> t^*$  }  $\Rightarrow$  each  $\leq d+p+1$

& base version = old base + all mods. at  $\leq t^*$  including  $t^*$

- use mods. to update reverses of  $\leq d+p$  pointers in base version of node' (node  $\rightarrow$  node')
- directly update reverses of pointers in mods. of node' (their time  $> t'$   $\Rightarrow$  don't care about node)



- potential  $\Phi = c \cdot \sum_{\text{node}} (\# \text{ mods. in second half of node})$    
 i.e. don't count first  $d+p+1$  mods.
- update creates 2 mods.  $\Rightarrow \Phi \uparrow 2c$    
 $\Rightarrow O(1)$  amortized cost ignoring splits
- split:  $\Phi \downarrow c(d+p+1)$  in node (empty second half)   
 &  $\Phi \uparrow c(d+p)$  to update reverses of base version of node'
- $\Rightarrow$  net  $\Phi \downarrow c$
- set  $c$  to the work of one split (excluding any recursive splits)
- $\Rightarrow$  amortized cost = 0

## De-amortization:

- partial:  $O(1)$  worst case [Brodal - NJC 1996]
- full: OPEN:  $O(1)$  worst case?



# Confluent persistence:



- after  $u$  confluent updates, can get size  $2^u$
- general transformation: [Fiat & Kaplan - J. Alg. 2003]
  - $d(v)$  = depth of version  $v$  in version DAG
  - $e(v) = 1 + \lg(\# \text{ paths from root to } v)$
  - overhead:  $\lg(\# \text{ updates}) + \underbrace{\max_v e(v)}_{\text{can be up to } u \dots}$  time & space


- still exponentially better than complete copy...

- lower bound:  $\sum e(v)$  bits of space [Fiat & Kaplan]
  - $\Rightarrow \Omega(e(v))$  for update if queries are free
  - construction makes  $\approx e(v)$  queries per update
- $O(\lg^3 \max_v e(v))$  update,  $O(\lg^2 \max_v e(v))$  query [Fiat & Kaplan]
- **OPEN**:  $O(1)$  or even  $O(\lg n)$  overhead per op.?

- disjoint transformation: [Collette, Iacono, Langerman - SODA 2012]
  - assume confluent ops. performed only on versions with no shared nodes
  - then  $O(\lg n)$  overhead possible

Idea: each node in subtree of version DAG

- only some of those versions modify node
- 3 types of versions:
  - node modified ~ easy
  - along path between mods.
  - below a leaf ~ hard
- fractional cascading [L3]  
& link-cut trees [L19]





# Functional: [Okasaki - book 2003]

path copying

- simple example: balanced BSTs
    - work top-down  $\Rightarrow$  no parent pointers
    - duplicate all changed nodes & ancestors before changing  $O(\lg n)$
- $\Rightarrow O(\lg n)/op.$

$\Rightarrow$  link-cut trees too [Demaine, Langerman, Price]

- e.g. deques with concat. in  $O(1)/op.$

double-ended queues [Kaplan, Okasaki, Tarjan - SICOMP 2000]

+ update & search in  $O(\lg n)/op.$

[Brodal, Makris, Tsihlias - ESA 2006]

- tries with local navigation & subtree copy/delete &  $O(1)$  fingers maintained to 'present'

[Demaine, Langerman, Price - Algorithmica 2010]

Think: Subversion

| method                    | finger move      |                  | modification             |
|---------------------------|------------------|------------------|--------------------------|
|                           | time             | space            | (time = space)           |
| path copying              | $\lg \Delta$     | $\emptyset$      | depth                    |
| 1 <sub>n</sub> functional | $\lg \Delta$     | $\lg \Delta$     | } local mods.<br>} cheap |
| 1 <sub>n</sub> confluent  | $\lg \lg \Delta$ | $\lg \lg \Delta$ |                          |
| 2 <sub>n</sub> functional | $\lg \Delta$     | $\emptyset$      | } globally<br>} balanced |
| 2 <sub>n</sub> confluent  | $\lg \lg \Delta$ | $\emptyset$      |                          |

