

1 Overview

This lecture is an interlude on time travel in data structures. We'll consider two kinds of time-travel (temporal) data structures, based on the two leading theories of time travel:

1. Branching-universe model a.k.a. persistence: You can change the past, but then you enter a different branch of the universe, and never to return to the old one. In some cases, though, branches might merge ("confluence").
2. Round-trip model a.k.a. retroactivity: You can change one thing in the past and then teleport back to the present and see what changed.

In data structures, persistence is much easier to achieve than retroactivity: in most cases, persistence can be done with tiny overhead, whereas retroactivity can require linear overhead. This is strong computational evidence that the second model of time travel is impossible, while the first model (and a model where you can go back, but not forward, in time) is conceivable computationally.

These data structural paradigms are also useful for more mundane applications like computational geometry, but we won't discuss those here.

2 Persistence

2.1 4 Levels of Persistence

We begin by describing the levels of desired persistence. With data structure persistence, we would like to keep all versions of the data structure available for updates and queries. Each persistence level, however will vary based on where updates are allowed and how branches and nodes are modified and created.

1. Partial Persistence – In this persistence model, we may query any previous version of the data structure, but we may only update the latest version. This implies a linear ordering among the versions.
2. Full Persistence – In this model, both updates and queries are allowed on any version of the data structure. The versions here form of a branching tree.
3. Confluent Persistence – In this model, we use combinators to combine input of more than one previous versions to output a new single version. Rather than a branching tree, combinations of versions induce a DAG(direct acyclic graph) structure on the version graph.

4. Functional Persistence – This model takes its name from functional programming where objects are immutable. The versions in this model are likewise immutable, so revisions do not alter the existing nodes in the data structure, but create new ones instead. Okasaki discusses these as well as other functional data structures in his book [?].

Each of the succeeding levels of persistence imply the preceding ones. That is, Functional Persistence \implies Confluent Persistence \implies Full Persistence \implies Partial Persistence. Functional implies confluent because we simply use the combinators to append a new combined version. Confluent implies full if we do not use combinators. Lastly, full implies partial if we only update the most recent version.




Figure 1: partial persistence: versions are linearly ordered.




Figure 2: full persistence: versions form a branching tree.




Figure 3: full persistence: versions form a DAG.

2.2 Partial Persistence

This result is due to Driscoll, Sarnak, Sleator, and Tarjan [?]. We work within the pointer machine model and require $O(1)$ in-degree per node, meaning that $\leq p = O(1)$ nodes point to any node. Each node stores some data and a constant number of pointers to children, reverse pointers to parents, and version modification data (in a “modification box”).

To maintain partial persistence, each node stores a reverse pointer to the parent node representing the most recent version of the data structure. A modification can be thought of as the tuple $(time, field, value)$, consisting of the time of the modification, the field being modified, and the

new value. We allow $\leq c + 1$ modifications in a node, c is the number of pointers per node, and is also $O(1)$.




Figure 4: nodes in the pointer machine.

To read files and check for nodes take $O(1)$ time. An update on a field at some time t can come across two cases:

- The node has space – We can simply add the modification $(t, field, value)$ to the modification box. All subsequent accesses of this modified node will check the modification box to override any initial data stored in the node.
- The node is full – We make a copy of the node, but using only the latest values. That is, we overwrite one of the nodes fields with the value that was stored in the modification box, and make the modification box of the new node empty. We propagate this change up to node's ancestors as follows: each ancestor makes a modification to change its child pointer to the newly created node. If that ancestor's modification box happens to be full, then we copy that node and propagate up. These changes propagate until we stop at the root.

Using Φ = the number of full nodes in the latest version. Since there are at most c pointers to point, there's at least one update spot in the modification box. Hence Φ can go down by 1. So we can prove constant time amortized updates. Further study by Brodal [?] has shown it to also be $O(1)$ in the worst case.

2.3 Full Persistence

This result is also due to [?]. We again assume a pointer machine with $\leq p$ incoming pointers per node. We can linearize the tree of versions via the in-order traversal marking the begin and end time of each version.

We can store the begin and end times in an order-maintenance data structure by Deitz and Sleator [?], discussed in lecture 17. The data structure can do the following two operations in $O(1)$ time

- insert time before or after a specified time




Figure 5: in-order traversal

- compare if time s precede time t . In our case, answer whether version v is an ancestor of version v' .

By the order-maintenance data structure, we can tell which modifications apply to the desired version. For each node, we store p back pointers, and allow up to $2(p + c + 1)$ modifications. As defined above, c is the pointers per node.

When a node is full, we split it into two nodes. But each is roughly half full (like B-tree node), then recursively update pointers and back pointers to this node. So even if all its c pointers move and all p back pointers move, the half full node - with $(p + c + 1)$ modifications - is still not full.

Again define Φ as the number of full nodes, after each split, it decreases at least one. We get $O(1)$ amortized cost. A related open question is

OPEN: Can we support $O(1)$ worst-case full persistence?

Each data structure node is represented by a linked list of nodes, and there's a second phase of the operation to update reverse pointers. Deitz developed a fully persistent array that can be achieved in $O(\lg \lg n) \times$ overhead in the word RAM model [?].

OPEN: Is there a matching lower bound for both full and partial persistence? This question may have been solved by Pătrascu et al. (unpublished).

2.4 Confluent Persistence

Confluent persistence has been explored in functional data structures [?]. Deques (double ended queues allowing stack and queue operations) with concatenation can be done in constant time per operation (Kaplan, Okasaki, and Tarjan [?]). We can create implicitly exponential deques in polynomial time by recursively concatenating a node with itself. The general transformation due to Fiat and Kaplan [?] is as follows:

- $d(v)$ = depth of node v in version DAG
- $e(v) = 1 + \lg(\text{number of paths root to } v)$
- overhead: $\lg(\text{number of updates}) + \max_v(e(v))$

- poor when $e(v) = 2^u$ where u is the number of updates. This is still exponentially better than the non-persistent model.




Figure 6: An example of $e(v)$ being exponential to the number of updates.

One example is Tries with $O(1)$ fingers to support local navigation and subtree copy and delete [?].

method	figure move		modification
	time	space	(time = space)
path copying	$\lg \Delta$	$O(1)$	depth
1. functional	$\lg \Delta$	$\lg \Delta$	$\lg \Delta$
1. confluent	$\lg \lg \Delta$	$\lg \lg \Delta$	$\lg \Delta$
2. functional	$\lg \Delta$	$O(1)$	$\lg n$
2. confluent	$\lg \lg \Delta$	$O(1)$	$\lg n$

The 1. functional and confluent persistence data structures are cheap with local modifications. The 2. data structures are globally balanced.

OPEN: Can we do better transformation with $O(1)$ fingers? separations? functional transformations?

OPEN: When can you do better? Lists with split and concatenate? Trees? General pointer machine? Array with cut and paste? Special DAGs? Others?

3 Retroactivity

Retroactive data structures store a data structure under a sequence of operations. We would like to be able to go back in time, change an operation, and then observe the effects of that change in the current state of the data structure. The induced timeline is linear. Much of this work is due to Demaine, Iacono, and Langerman [?].

The allowed operations are

- $\text{Insert}(t, x)$ – Retroactively do operation x at time t
- $\text{Delete}(t)$ – Retroactively undo operation at time t
- $\text{Query}(t, x)$ – Do query x at time t

We define partial retroactivity as allowing queries at the present time and full retroactivity as allowing queries at any time. Some cases of partial retroactivity are easy to implement. If updates are commutative: $x \circ y = y \circ x$, then we can support retroactive insertion of operations at no additional asymptotic cost (implement $\text{Insert}(t, x)$ by executing x at the present time). If updates, in addition to being commutative, are also invertible: $x \circ x^{-1} = \text{NOP}$, then we can support partial retroactivity at no additional asymptotic cost ($\text{Insert}(t, x)$ by executing x at the present time and $\text{Delete}(t)$ by executing x^{-1} at the present time where x^{-1} is the inverse of the operation at time t).

3.1 The Rollback Method

There are a few general transformations that we can prove bounds for. One is the rollback method, in which we perform a retroactive operation at r time units in the past. We can do this with a factor of r overhead by keeping a log of all updates done to the DS such that every change can be reversed. The rollback method needs an $\Omega(r)$ lower bound. To see this, we examine a data structure that maintains two values, X and Y , both initialized to 0. We can perform the following operations on our data structure:

- $\text{set}X(x)$ – Sets $X \leftarrow x$
- $\text{add}Y(\Delta)$ – Sets $Y \leftarrow Y + \Delta$
- $\text{mult}XY()$ – Sets $Y \leftarrow X \cdot Y$
- $\text{query}()$ – Returns Y

Consider the following sequence of operations: $\text{add}Y(a_n)$, $\text{mult}XY()$, $\text{add}Y(a_{n-1})$, $\text{mult}XY()$, ..., $\text{add}Y(a_0)$. This is Horner’s rule for evaluating the polynomial $p(x) = \sum_{i=0}^n a_i x^i$. Now suppose we perform $\text{Insert}(t = 0, \text{set}X(x_0))$ to change the x value of the evaluated polynomial x_0 . Frandsena, Hansen, and Miltersen in 2001 showed that evaluating a polynomial of degree n requires $\Omega(n)$ time over any field, independent of any preprocessing of the a_i s. This holds in the “history-independent algebraic decision tree” model, which implies the same result for the integer RAM and generalized real RAM models [?]. This is somewhat disappointing result because it says that in the retroactive model, we can’t do any persistence maintenance that’s better than just going back in time, performing the new operation, and then re-executing all of the operations in our history past that point.

In the cell-probe model, Frandsena et al. also proved a lower bound of $\Omega(\sqrt{r/\lg r})$. They had a data structure that maintained n words and supported arithmetic updates $(+/\cdot)$. Computing a fast fourier transform takes $O(n \lg n)$ time, but changing one weights w_i of the FFT needs $\Omega(\sqrt{n})$ time, from which we derive the $\Omega(\sqrt{r/\lg r})$ lower bound. An open question is whether the tightest cell-probe lower bound is $\Omega(\sqrt{r/\text{poly} \lg r})$.

3.2 Priority Queues

Let us turn our attention to partially retroactive priority queues. The defining features of priority queues is the `delete-min()` operation, which makes the set of operations on priority queues non-commutative. We can plot the status of our data structure in the plane. The x -axis represents time and y -axis represents key value. Every `insert(t, k)` operation creates a horizontal ray that starts at point (t, k) and shoots to the right. Every `delete-min()` operation creates a vertical ray that starts at $(t, -\infty)$ and shoots upwards, stopping at the horizontal ray of the element it deletes. It turns the horizontal ray into a line segment with endpoints (t, k) and (d_k, k) , where d_k is the time of key k 's deletion. This creates nonintersecting upsidedown “L” shapes, where each L corresponds to an `insert` and the `delete-min()` that deletes it. Refer to Figure ?? for an animation.

Let Q_0 be the current state of our priority queue and Q_t be its state at time t . We call time t a *bridge* if $Q_t \subseteq Q_0$. There are four combinations of retroactive operations:

1. `Insert($t, \text{“insert}(k)\text{”}$)` – Insert key k into Q_t . The resulting element we insert into Q_0 is the largest element that was deleted after time t . See Figure ??.
2. `Delete($t, \text{“delete-min()”}$)` – Undo the `delete-min` at time t . This is identical to re-inserting the element that was being deleted at the time of deletion (i.e. nullify the upwards `delete-min()` arrow by inserting the appropriate key right at that time). Thus, it is the same as the above case and we insert into Q_0 the largest element that was deleted after time t .
3. `Insert($t, \text{“delete-min()”}$)` – Delete the minimum key at time t . The element we delete from Q_0 is the minimum value of $Q_{t'}$, where t' is the first bridge after t . This essentially pushes the bridge forward in time. See Figure ??.
4. `Delete($t, \text{“insert}(k)\text{”}$)` – Undo the insertion of key k at time t . If $k \in Q_0$ we remove it from there. If not, then we again delete the minimum value of $Q_{t'}$ where t' is the first bridge after t . The idea is that since $k \notin Q_0$, it didn't make it to its next bridge. Therefore, removing the insertion of that number will cascade up the deletes before that bridge, so the minimum element from the that bridge will get removed by the last cascaded delete.

We can perform all of these operations in $O(\lg m)$ worst-case, where m is the total number of updates, present or retroactive, performed on the priority queue. We first store a balanced BST insertions keyed on insertion time. We augment each node of this tree with the max key $k' \notin Q_0$ over every node's subtree. This lets us find the maximum key among all elements inserted after a time t' but not in Q_0 in $O(\lg m)$ time. We can also find the minimum key among all elements inserted before a time t' and still in Q_0 in the same runtime if we maintain in every node the minimum of all keys in its subtree still in Q_0 . These are useful for operations 1 and 2, and 3 and 4, respectively.

To find the last bridge before t or the first bridge after t , we maintain a list of updates which we store in a modified (a, b) -tree developed by Fleischer [?]. If we assign a weight of 0 to `insert(k)` for $k \in Q_0$, +1 to `insert(k)` for $k \notin Q_0$, and -1 for `delete-min()`, every bridge in the tree corresponds to a prefix sum of 0. This allows us to find bridges in $O(\lg m)$ time, which we use for every operation.

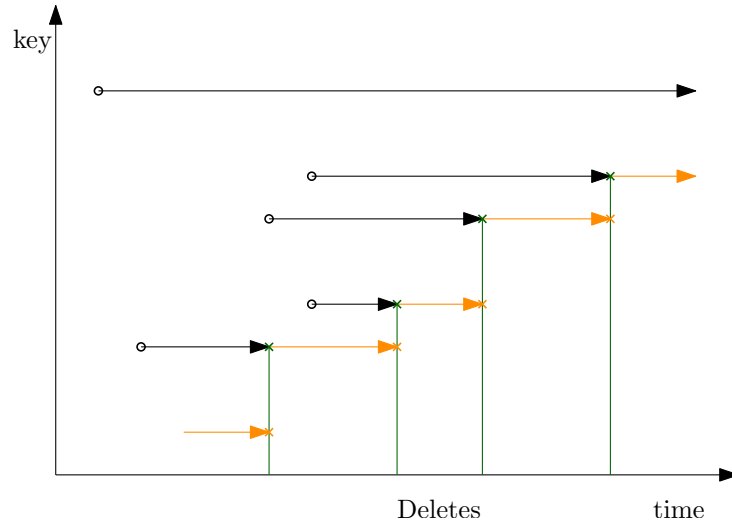


Figure 7: Priority queues.




Figure 8: Successor.

3.3 Other Data Structures

- queue: partial retroactivity is $O(1)$, and full retroactivity is $O(\lg m)$.
- deque: full retroactivity is $O(\lg n)$.
- union-find (incremented connectivity): full retroactivity is $O(\lg m)$.
- priority queue: partial retroactivity is $O(1)$, and full retroactivity is $O(\lg m)$.
- successor: $O(\lg m)$ for partial retroactivity. The easy approach to full retroactivity takes $O(\lg^2 m)$. Giora and Kplan [?] can achieve $O(\lg m)$ time for full retroactivity. Finding the successor is equivalent to optimal dynamic vertical ray shooting among horizontal line segments (see the figure below).

OPEN: Can we solve optimal dynamic vertical ray shooting among general (not necessary horizontal) line segments?

References

- [1] M. Fredman, J. Komlós, E. Szemerédi, *Storing a Sparse Table with $O(1)$ Worst Case Access Time*, Journal of the ACM, 31(3):538-544, 1984.
- [2] Gerth Stlting Brodal: *Partially Persistent Data Structures of Bounded Degree with Constant Update Time*. Nord. J. Comput. 3(3): 238-255 (1996)
- [3] Erik D. Demaine, John Iacono, Stefan Langerman: *Retroactive data structures*. SODA 2004: 281-290
- [4] Erik D. Demaine, Stefan Langerman, and Eric Price: *Confluently Persistent Tries for Efficient Version Control*. Algorithmica (2008).
- [5] Paul F. Dietz, Daniel Dominic Sleator: *Two Algorithms for Maintaining Order in a List* STOC 1987: 365-372
- [6] Paul F. Dietz: *Fully Persistent Arrays (Extended Array)*. WADS 1989: 67-74
- [7] James R. Driscoll, Neil Sarnak, Daniel Dominic Sleator, Robert Endre Tarjan: *Making Data Structures Persistent*. J. Comput. Syst. Sci. 38(1): 86-124 (1989)
- [8] Yoav Giora, Haim Kaplan: *Optimal dynamic vertical ray shooting in rectilinear planar subdivisions* ACM Transactions on Algorithms 5(3) (2009)
- [9] Haim Kaplan, Chris Okasaki, Robert Endre Tarjan: *Simple Confluently Persistent Catenable Lists*. SIAM J. Comput. 30(3): 965-977 (2000)
- [10] Amos Fiat, Haim Kaplan: *Making data structures confluently persistent*. J. Algorithms 48(1): 16-58 (2003)
- [11] Chris Okasaki: *Purely Functional Data Structures*. New York: Cambridge University Press, 2003.
- [12] Gudmund Skovbjerg Frandsen, Gudmund Skovbjerg Frandsen, Peter Bro Miltersen: *Lower Bounds for Dynamic Algebraic Problems*. Information and Computation 171(2): 333-349 (2001)
- [13] Rudolf Fleischer: *A Simple Balanced Search Tree with $O(1)$ Worst-Case Update Time*. Int. J. Found. Comput. Sci. 7(2): 137-150 (1996)