



Amortized Analysis of Splay Trees via a Lax Homomorphism

Lukas Kebuladze

Carnegie Mellon University

On-Paper Amortized Analysis (Potential Method)

Amortized analysis, instead of analyzing individual operations like in standard asymptotic analysis, considers the total cost of a sequence operations as follows.

- Define *potential function* $\Phi : D \rightarrow \mathbb{R}$ where D is the set of data structure states.
- $\forall d \in D, \Phi(d)$ represents the *potential* (difference between real and imagined cost).
- Given $\delta_o : D \rightarrow D$ that implements an operation, then for any arbitrary $d \in D$, the *amortized cost* of δ_o on state d is defined as

$$\sigma_{\S} = \delta_{\S}(d) + \Phi(\delta_o(d)) - \Phi(d) \quad (1)$$

where $\delta_{\S} : D \rightarrow \mathbb{R}$ is the true cost of the operation on state d . Letting $d^{(i)} = \delta_o^{(i)}(d)$ and $\sigma_{\S}^{(i)} = \delta_{\S}(d^{(i)}) + \Phi(d^{(i+1)}) - \Phi(d^{(i)})$, the following holds by telescoping sums:

$$\sum_{i=0}^{m-1} \delta_{\S}(d^{(i)}) = \Phi(d^{(0)}) - \Phi(d^{(m)}) + \sum_{i=0}^{m-1} \sigma_{\S}^{(i)} \quad (2)$$

Key idea: If $\Phi(d^{(0)}) - \Phi(d^{(m)}) \leq 0$, then $\sum_{i=0}^{m-1} \delta_{\S}(d^{(i)}) \leq \sum_{i=0}^{m-1} \sigma_{\S}^{(i)}$ according to Equation 2 which bounds the true cost of a sequence of m operations in terms of amortized cost.

Cost-Aware Logical Framework

What it is: The directed effectful cost-aware logical framework (**decalf**) is a *directed* dependent call-by-push-value language specialized for cost analysis [3].

- Call-by-push-value (cbpv) distinguishes pure values and effectful computations, and employs a pair of adjoint functors $F \dashv U$ to mediate between them.
- Of particular importance in cbpv are **bind**, which sequences computations of F types, and **ret**, which returns values as effectless computations, which together satisfy the standard β and η equations.

$$\frac{F\text{-Intro} \quad a : A}{\text{ret}(a) : F(A)} \quad \frac{F\text{-Elim} \quad e : F(A) \quad f : A \rightarrow X}{\text{bind}(e; f) : X}$$

Notationally we write $A \rightarrow B := U(A \rightarrow F(B))$ for the (effectful) Kleisli functions.

Accounting of cost: Cost is reified in decalf as a computation $\text{step}^c(e)$ that records c units of cost (over a cost monoid \mathbb{C}) before running the computation e .

Program inequality at each type: $e \leq_X e'$ at computation type X indicates that the cost of computation e is bounded by that of e' and they return the same result. This provides an integrated way to perform simultaneous cost and correctness verification.

$$\begin{array}{l} \text{step}^0(e) = e \\ \text{step}^{c_1}(\text{step}^{c_2}(e)) = \text{step}^{c_1+c_2}(e) \\ \text{bind}(\text{step}^c(e); f) = \text{step}^c(\text{bind}(e; f)) \end{array} \quad \begin{array}{l} c \leq_{\mathbb{C}} c' \rightarrow \text{step}^c(e) \leq_X \text{step}^{c'}(e) \\ e \leq_{F(A)} e' \rightarrow ((x : A) \rightarrow f(x) \leq_X f'(x)) \\ \rightarrow \text{bind}(e; f) \leq_X \text{bind}(e'; f') \end{array}$$

Figure 1. Monoid structure of cost and bind commutativity with cost

Figure 2. Monotonicity at **step** and **bind**

Splay Trees

A *splay tree* is a self-adjusting binary search tree (BST) developed by Sleator and Tarjan [4] where **every time a node is accessed, it is moved to the root of the tree**.

It moves the node to the root via the operation $\text{splay} : \mathbb{N} \times \text{tree} \rightarrow \text{tree}$, which finds key $k : \mathbb{N}$ in tree $t : \text{tree}$, and then rotates k to the root of t via a sequence of single or double rotations, called *splay-steps*, while preserving in-order traversal.

Goal: show that for an n -node tree, **splay** has $O(\log n)$ amortized cost and m **splay** operations on an n -node tree has actual cost $O(m \log n + n \log n)$.

Understanding the Splay Operation

The efficiency of **splay** comes from the *splay-steps*, which are divided into six cases, three of which are symmetric. Three cases are depicted below (left) with node x being splayed and an example of **splay** operating on a tree (right).

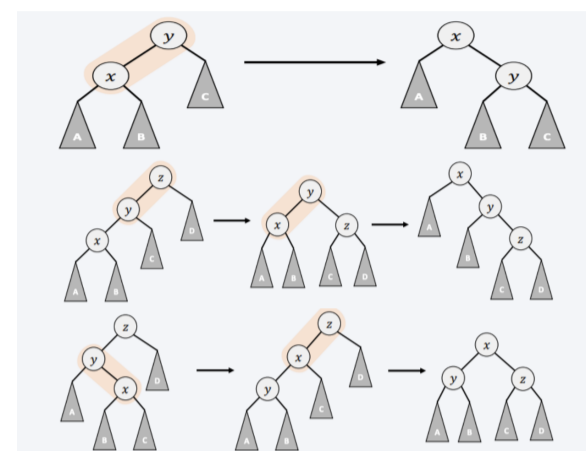


Figure 3. From top to bottom: zig, zig-zig, zig-zag




Figure 4. Tree becomes more balanced after splaying

How to Mechanize Amortized Analysis

Given some amortized data structure \mathbf{D} with operation δ charging the real cost and specification data structure \mathbf{S} with operation σ that charges the amortized cost of δ . In **decalf** embedded in Agda, do as follows.

- Implement (\mathbf{D}, δ) and (\mathbf{S}, σ) annotated with **step**, ascribing to some algebraic signature.
- Define a lax homomorphism $\varphi : \mathbf{D} \rightarrow \mathbf{S}$ that charges potential function $\Phi : \mathbf{D} \rightarrow \mathbb{R}$.
- Show that the single square commutes which proves the amortized cost bound of a single application of δ .
- Show that the m -fold square commutes which proves the actual cost bound of m applications of δ .

Why Use a Lax Homomorphism?

This approach is prior work outlined in Grodin and Harper [1] and its application to this case is inspired by Grodin et al. [2]. The contribution here is applying this technique to splay trees.

Why Homomorphism? If we rearrange Equation 1 by moving $\Phi(d)$ to the other side, we get this:

$$\delta_{\S}(d) + \Phi(\delta_o(d)) = \Phi(d) + \sigma_{\S}$$

which is exactly the form of a homomorphism φ which charges cost Φ (i.e. a *single commuting square*). Since such a homomorphism also proves observational equivalence of states, the homomorphism φ simultaneously defines the relation between the two states and the potential function.

Once this is proved, Equation 2 can be easily achieved by a m -fold composition of the single commuting square

Why Lax? Complex amortized data structures, such as splay trees, have imprecise amortized analysis, since the amortized cost is an *upper bound* on the actual cost. In other words, Equation 1 needs to be turned into an inequality like so:

$$\delta_{\S}(d) + \Phi(\delta_o(d)) \leq \Phi(d) + \sigma_{\S}$$

which is achieved via a lax homomorphism.

Mechanization Definitions

First, the signature, amortizing implementation, and specification are defined.

$$\begin{array}{lll} \text{record BST where} & \mathbf{ST} : \text{BST} & \mathbf{LT} : \text{BST} \\ \text{T} : \text{tp}^+ & \mathbf{ST.T} = \text{tree} & \mathbf{LT.T} = \text{list} \\ \text{find} : \mathbb{N} \times \text{T} \rightarrow \text{T} & \mathbf{ST.find} = \text{splay} & \mathbf{LT.find}(k, l) = \text{step}^{3\lceil \log_2 l \rceil + 1}(\text{ret}(l)) \end{array}$$

Figure 5. Signature

Figure 6. Amortizing Impl

Figure 7. Specification

For demonstration purposes, this signature is slightly simplified. The complete type of **find** is $\mathbb{N} \times \text{T} \rightarrow \text{bool} \times \text{T}$ where the boolean indicates whether the key k is found in tree t . This complete type rules out trivial implementations and allows for stronger proofs of correctness. Additionally, the signature includes a constructor method, making the signature and abstraction for static sets.

Mechanizing the Lax Homomorphism

Then, the lax homomorphism φ is defined, which charges the potential Φ and returns the in-order traversal (**inord**) representation of the tree.

$$\begin{array}{ll} \Phi : \mathbf{ST.T} \rightarrow \mathbb{N} & \varphi : \mathbf{ST.T} \rightarrow \mathbf{LT.T} \\ \Phi(\text{leaf}) = 0 & \varphi(t) = \text{step}^{\Phi(t)}(\text{ret}(\text{inord}(t))) \\ \Phi(\text{node } l \ x \ r) = \Phi(l) + \lceil \log_2 \lceil (\text{node } l \ x \ r) \rceil \rceil + \Phi(r) & \end{array}$$

Note that Φ is equivalent to the potential function used in the original analysis by Sleator and Tarjan [4].

Single Commuting Square

Through structural induction on the depth of the target key in the tree and reasoning about splay trees similar to the original amortized analysis by Sleator and Tarjan, the following is mechanized:

Generalized Access Lemma

For all keys k and splay trees t it follows that

$$\text{bind}(\mathbf{ST.find}_k(t); \varphi) \leq_{F(\mathbf{LT.T})} \text{bind}(\varphi(t); \mathbf{LT.find}_k).$$

This is equivalent to the following commutative diagram over the Kleisli category:

$$\begin{array}{ccc} \mathbf{ST.T} & \xrightarrow{\mathbf{ST.find}_k} & \mathbf{ST.T} \\ \varphi \downarrow & \lrcorner & \downarrow \varphi \\ \mathbf{LT.T} & \xrightarrow{\mathbf{LT.find}_k} & \mathbf{LT.T} \end{array}$$

This lemma and corresponding diagram is the precise condition for φ to be a valid lax homomorphism.

Composed Commuting Square

To state the final theorem capturing the cost of a sequence of m **splay** operations, then $\text{IsBounded}(e, c) := e; \text{ret}(\star) \leq_{F(1)} \text{step}^c(\text{ret}(\star))$ is used to focus only on the cost to get the following which is proved in **decalf**.

Balance Theorem

Let k_1, k_2, \dots, k_m be arbitrary keys where m is the number of **finds** performed. Let n be the number of nodes in the input splay tree t . Then,

$$\text{IsBounded}(\mathbf{ST.find}_{k_i}^m(t)) (m(3\lceil \log_2 n \rceil + 1) + n\lceil \log_2 n \rceil)$$

where $\mathbf{ST.find}_{k_i}^m$ is the m -fold application of **find** on keys k_1, k_2, \dots, k_m .

Similarly, this is equivalent to the following composed commutative diagram:

$$\begin{array}{ccccccc} \mathbf{ST.T} & \xrightarrow{\mathbf{ST.find}_{k_1}} & \mathbf{ST.T} & \longrightarrow & \dots & \longrightarrow & \mathbf{ST.T} & \xrightarrow{\mathbf{ST.find}_{k_m}} & \mathbf{ST.T} \\ \varphi \downarrow & \lrcorner & \downarrow \varphi & & & & \downarrow \varphi & \lrcorner & \downarrow \varphi \\ \mathbf{LT.T} & \xrightarrow{\mathbf{LT.find}_{k_1}} & \mathbf{LT.T} & \longrightarrow & \dots & \longrightarrow & \mathbf{LT.T} & \xrightarrow{\mathbf{LT.find}_{k_m}} & \mathbf{LT.T} \end{array}$$

References

- Harrison Grodin and Robert Harper. Amortized Analysis via Coalgebra. *Electronic Notes in Theoretical Informatics and Computer Science*, Volume 4 - Proceedings of MFPS XL, December 2024.
- Harrison Grodin, Running Li, and Robert Harper. Abstraction functions as types. *Proc. ACM Program. Lang.*, 10(POPL), January 2026.
- Harrison Grodin, Yue Niu, Jonathan Sterling, and Robert Harper. Decalf: A Directed, Effectful Cost-Aware Logical Framework. *Proceedings of the ACM on Programming Languages*, 8(POPL):10:273–10:301, January 2024.
- Daniel Dominic Sleator and Robert Endre Tarjan. Self-adjusting binary search trees. *J. ACM*, 32(3):652–686, July 1985.