

Reasoning about B+ Trees with Operational Semantics and Separation Logic

Hayo Thielecke
(joint work with Alan Sexton)

University of Birmingham, UK

Introduction

B+ tree

- ▶ Balanced ordered search tree, $O(\log n)$ operations
- ▶ Widely used for efficient search in databases

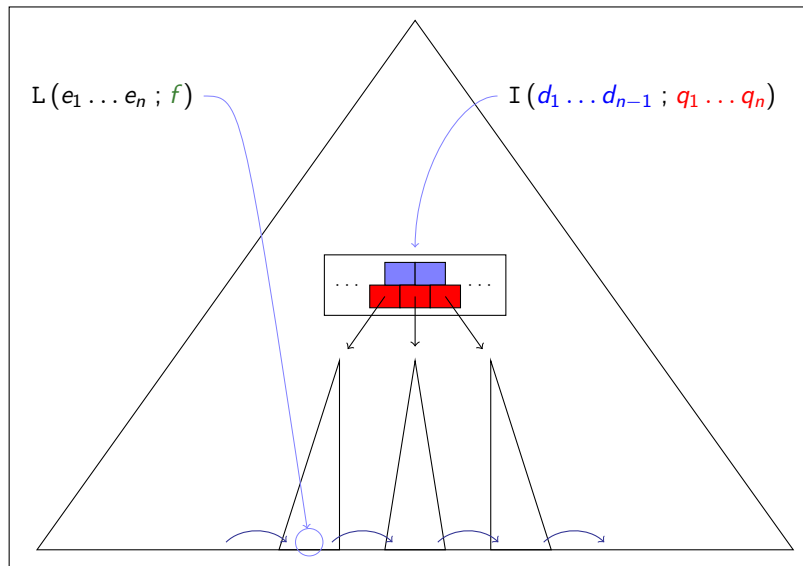
Operational semantics: Abstract Machine

- ▶ Formalizes operations on B+ trees, abstractly
- ▶ In the style of SECD machine or Krivine's machine

Separation logic

- ▶ Very expressive for inductive data structures with pointers
- ▶ Here: ordinary predicate logic extended with $*$ and \mapsto
- ▶ No Hoare logic $\{P\} c \{Q\}$

B+ tree visually



Operational semantics: abstract machine

Landin's SECD machine (1964)

Configurations: Sack, Environment, Control, Dump

$$\langle s, e, c, d \rangle \rightsquigarrow \langle s', e', c', d' \rangle$$

Transitions by pattern matching (e.g., stack or control empty).

Our machine: insertion transitions

$$\langle C, r, \pi, \sigma \rangle \rightsquigarrow \dots$$

C : command, e.g., $\text{Insert}(e)$, S , $D(k, q)$

r : root of current subtree

π : stack, used for tree traversal

σ : store, holds nodes of the tree, \approx heap in Sep Logic

Separation logic connectives

$\sigma \models p \mapsto N$ iff $\sigma = \{p \mapsto N\}$ for a node N (leaf or internal)

$\sigma \models Q_0 * Q_1$ iff $\sigma = \sigma_0 \cup \sigma_1$ where $\sigma_0 \cap \sigma_1 = \emptyset$

and $\sigma_0 \models Q_0$ and $\sigma_1 \models Q_1$

$\sigma \models Q_0 \wedge Q_1$ iff $\sigma \models Q_0$ and $\sigma \models Q_1$

$\sigma \models Q_0 \vee Q_1$ iff $\sigma \models Q_0$ or $\sigma \models Q_1$

$\sigma \models \text{true}$ iff σ is any store

$\sigma \models \text{emp}$ iff $\text{dom}(\sigma) = \emptyset$

Local reasoning: **footprint** and **frame**:

$$\frac{\{P\} c \{Q\}}{\{P * R\} c \{Q * R\}}$$

B+ trees in Separation Logic

Predicate describes data structure

$\sigma \models \text{Btree}_h(r, S, a, z, n)$

- ▶ height h
- ▶ r points to root node of tree
- ▶ set of entries S
- ▶ first leaf at address a , last leaf node points to z
- ▶ size of root is at least n

Induction on height: Leaf case

$$\begin{aligned} \text{Btree}_1(r, S, a, z, n) &\iff \exists e_1, \dots, e_n. n \leq \text{MaxN} \\ &\wedge r \mapsto L(e_1 \dots e_n; z) \\ &\wedge S = \{e_1, \dots, e_n\} \wedge a = r \wedge e_1 \sqsubset \dots \sqsubset e_n \end{aligned}$$

B+ tree continued: internal nodes

$$\begin{aligned}
 \text{Btree}_{h+1}(r, S, a, z, n) &\iff \exists d_1, \dots, d_{n-1}, q_1 \dots q_n, m_1, \dots, m_n. \\
 &\wedge (r \mapsto \mathbf{I}(d_1 \dots d_{n-1}; q_1 \dots q_n)) \\
 &\quad * \text{Btree}_h(q_1, S_1, a_1, a_2, m_1) \\
 &\quad * \text{Btree}_h(q_2, S_2, a_2, a_3, m_2) \\
 &\quad * \dots \\
 &\quad * \text{Btree}_h(q_n, S_n, a_n, a_{n+1}, m_n) \\
 &\wedge a_1 = a \wedge a_{n+1} = z \\
 &\wedge S = S_1 \cup \dots \cup S_n \\
 &\wedge (\forall j. 1 < j < n - 1 \Rightarrow d_j \sqsubseteq S_j \sqsubseteq d_{j+1}) \\
 &\wedge (\forall j. 1 < j \leq n \Rightarrow \lceil \text{MaxN}/2 \rceil \leq m_j) \\
 &\wedge (S_1 \sqsubseteq d_1) \\
 &\wedge (d_{n-1} \sqsubseteq S_n)
 \end{aligned}$$

B+ tree insertion

- ▶ Invariant: **balanced**
- ▶ **Insertion** may **split** a node.
- ▶ **Parent node** may then split as well.
- ▶ Splitting ripples **up the tree recursively**.
- ▶ Insertion produces either **a subtree** or **two subtrees**.
- ▶ In the operational semantics: “**single**” or “**double**” commands

$$\langle \mathbf{S}, r, \pi, \sigma \rangle$$

$$\langle \mathbf{D}(k, q), r, \pi, \sigma \rangle$$

Insert Rules — Descending the tree and splitting leaf

$$\langle \text{Insert}(a), r, \pi, \sigma \rangle \rightsquigarrow \langle \text{Insert}(a), \mathbf{p}_i, (r, i) :: \pi, \sigma \rangle$$

if $\sigma(r) = \mathbf{I}(\mathbf{d}; \mathbf{p})$
where $i = \text{first}(\mathbf{d}, \lambda x. x > \text{key}(a))$

$$\langle \text{Insert}(a), r, \pi, \sigma \rangle \rightsquigarrow \langle \mathbf{D}(k, q), r, \pi, \sigma \left[\begin{array}{l} r \mapsto \mathbf{L}(\mathbf{e}'; q), \\ q \mapsto \mathbf{L}(\mathbf{e}''; f) \end{array} \right] \rangle$$

where $\sigma(r) = \mathbf{L}(\mathbf{e}; f)$
and $i = \text{first}(\mathbf{e}, \lambda x. \text{key}(x) \geq \text{key}(a))$
and $\langle \mathbf{e}', k, \mathbf{e}'' \rangle = \text{splitL}(i, a, \mathbf{e})$
and $q \notin \text{dom}(\sigma)$

Insert Rules — Returning up the tree with S or D(k, q)

$$\langle \mathbf{S}, r, (t, i) :: \pi, \sigma \rangle \rightsquigarrow \langle \mathbf{S}, t, \pi, \sigma \rangle$$

$$\langle \mathbf{D}(k, q), r, (t, i) :: \pi, \sigma \rangle \rightsquigarrow \langle \mathbf{S}, t, \pi, \sigma[t \mapsto \mathbf{I}(\mathbf{d}' ; \mathbf{p}')] \rangle$$

if $|\mathbf{p}| < \mathit{MaxN}$

where $\sigma(t) = \mathbf{I}(\mathbf{d} ; \mathbf{p})$

and $\mathbf{d}' = \mathit{ins}(k, i, \mathbf{d})$

and $\mathbf{p}' = \mathit{ins}(q, i + 1, \mathbf{p})$

$$\langle \mathbf{D}(k, q), r, (t, i) :: \pi, \sigma \rangle \rightsquigarrow \langle \mathbf{D}(k', q'), t, \pi, \sigma \left[\begin{array}{l} t \mapsto \mathbf{I}(\mathbf{d}' ; \mathbf{p}'), \\ q' \mapsto \mathbf{I}(\mathbf{d}'' ; \mathbf{p}'') \end{array} \right] \rangle$$

where $\sigma(t) = \mathbf{I}(\mathbf{d} ; \mathbf{p})$

and $\langle \mathbf{d}', \mathbf{p}', k', \mathbf{d}'', \mathbf{p}'' \rangle = \mathit{splitI}(i, k, q, \mathbf{d}, \mathbf{p})$

and $q' \notin \mathit{dom}(\sigma)$

Correctness of insertion Idea: *footprint* * *frame* for the same stack π .

Lemma

Assume $\sigma \models \text{Btree}_h(r, S, a, z, n) * R$. Then either:

1. $\langle \text{Insert}(e), r, \pi, \sigma \rangle \rightsquigarrow \dots \rightsquigarrow \langle S, r, \pi, \sigma' \rangle$ and

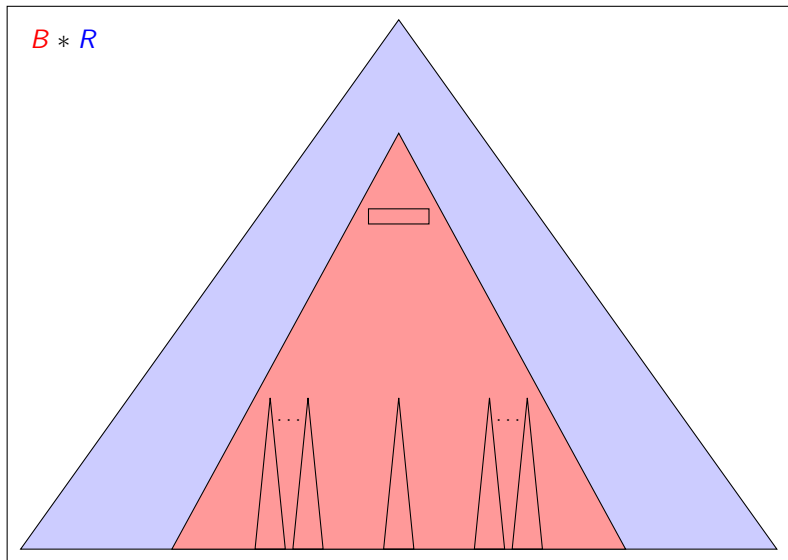
$$\sigma' \models \text{Btree}_h(r, S + e, a, z, n) * R$$

2. $\langle \text{Insert}(e), r, \pi, \sigma \rangle \rightsquigarrow \dots \rightsquigarrow \langle D(k, q), r, \pi, \sigma' \rangle$ and


$$\sigma' \models \text{Btree}_h(r, S_r, a, b, n') * \text{Btree}_h(q, S_q, b, z, n'') * R$$

Moreover, $S_r \cup S_q = S + e$ and $S_r \sqsubset k \sqsubseteq S_q$.

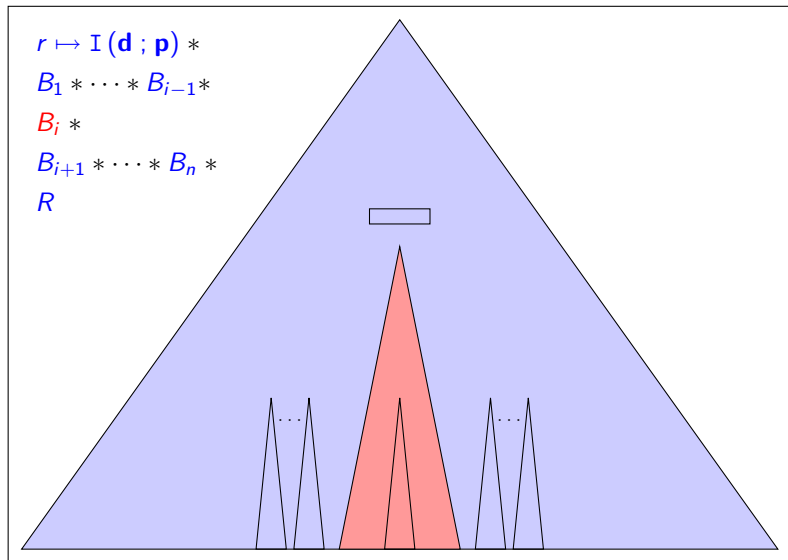
Proving Correctness of Insertion



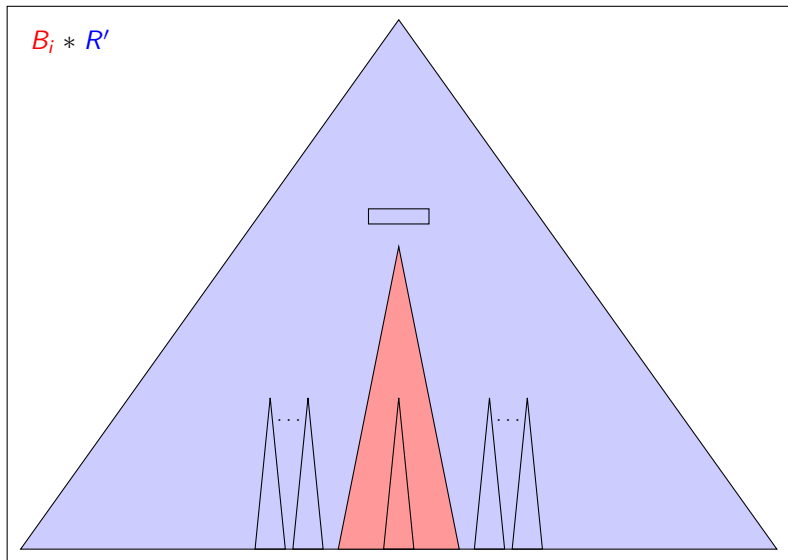
Proving Correctness of Insertion



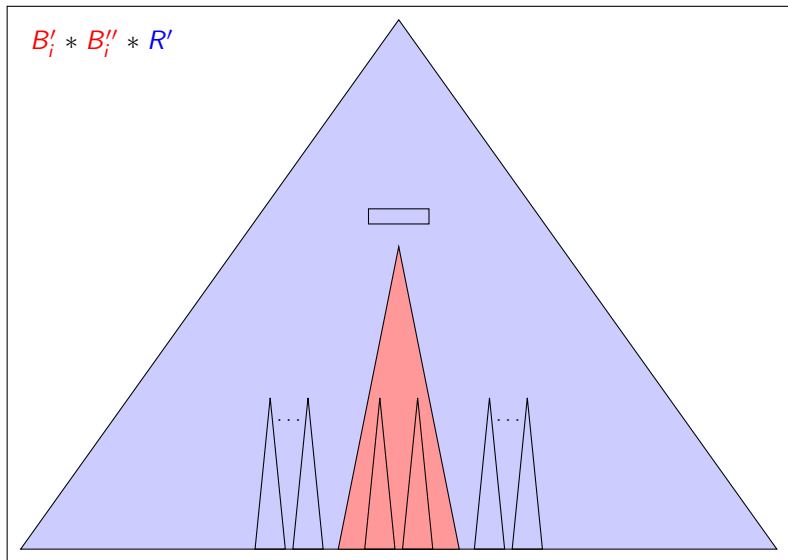
Proving Correctness of Insertion




Proving Correctness of Insertion



Proving Correctness of Insertion



Proving Correctness of Insertion



Finding a list of entries for a range query

- ▶ Find the list of entries with $\text{key} \geq k$.
- ▶ Need to find the first such element at the leaf level
- ▶ Use **fringe links** of the leaf nodes
- ▶ Machine descends the tree (no need to keep a stack).

$$\langle \text{Find}(k), r, \sigma \rangle \rightsquigarrow \langle \text{Find}(k), \mathbf{p}_i, \sigma \rangle$$

if $\sigma(r) = \mathbf{I}(\mathbf{d}; \mathbf{p})$
where $i = \text{first}(\mathbf{d}, \lambda x. x > k)$

$$\langle \text{Find}(k), r, \sigma \rangle \rightsquigarrow \langle \text{Ret}(r, i), \sigma \rangle$$

where $\sigma(r) = \mathbf{L}(\mathbf{e}; f)$
and $i = \text{first}(\mathbf{e}, \lambda x. \text{key}(x) \geq k)$

Correctness of find

- ▶ Start from a B+ tree with its **implicit fringe list**
- ▶ Result of find is covered by a **list predicate**
- ▶ Rest of the tree is swept into the catch-all predicate **true**.
- ▶ **Gerrymander** the existing store to reveal the result list.

Theorem


If $\sigma \models \exists h, a. \text{Btree}_h(r, S, a, \text{null}, n)$, then

$$\langle \text{Find}(k), r, \sigma \rangle \rightsquigarrow \dots \rightsquigarrow \langle \text{Ret}(q, i), \sigma \rangle$$

for some q and i with

$$\sigma \models \underbrace{\text{FList}(q, i, S \uparrow k)}_{\text{list of results}} \quad * \quad \underbrace{\text{true}}_{\text{everything else}}$$

Correctness of find Appending the fringe list onto an accumulator.



Lemma

Suppose

$$\sigma \models \text{Btree}_h(r, S_r, a, z, n) * \text{FList}(z, 1, S_z)$$

Then $\sigma \models \text{FList}(a, 1, S_r \cup S_z) * \text{true}$.

Correctness of find: induction

A technical lemma combines:

- ▶ current subtree
- ▶ partial result list
- ▶ everything else, by way of true

Lemma

Let

$$\sigma \models \text{Btree}_h(p, S_p, a, z, n) * \text{FList}(z, 1, S_z) * \text{true}$$

Then we have a sequence of h transitions

$$\langle \text{Find}(k), r, \sigma \rangle \rightsquigarrow \dots \rightsquigarrow \langle \text{Ret}(q, i), \sigma \rangle$$

for some q and i with $\sigma \models \text{FList}(q, i, (S_r \uparrow k) \cup S_z) * \text{true}$.

Conclusions

Summary

- ▶ Programming semantics technology to address B+ trees
- ▶ Separation logic \mapsto and $*$ for invariants and local reasoning
- ▶ Abstract machines from functional programming
- ▶ Translation to CAML prototype is straightforward

Further work

- ▶ Deletion will be covered in the full version of the paper
- ▶ High-level primitives, like Gardner's context logic for XML
- ▶ Trees with "holes", holey Brick and BV tree: a job for \rightarrow^* ?
- ▶ Concurrent B trees

The End

Thank you for your attention.