

Flow Interfaces

Compositional Abstractions for Concurrent Data Structures

Siddharth Krishna, Dennis Shasha, and Thomas Wies


Motivation

Verifying concurrent data structures by only reasoning about the small region modified by each thread (**compositional** reasoning).

Challenges

- Unbounded sharing and complex overlays
- Data invariants depend on global shape

Examples: Harris' non-blocking list (below), B-link trees



Current approaches

- Separation logic (SL) based logics
- *Inductive predicates* to describe shape and data properties
- Example: list segments

$$ls(x, y) := (x = y \wedge emp) \vee (\exists z. x \mapsto z * ls(z, y))$$

- **Problem 1:** definition tied to traversal that visits every node exactly once
 - How do we describe Harris' list?

- **Problem 2:** predicates and lemmas are data-structure-specific

- List composition:

$$ls(x, y) * ls(y, z) \Rightarrow ls(x, z)$$

- Sorted list segment with upper and lower bounds: $sls(x, y, l, u)$

- Different composition:

$$sls(x, y, l, v) * sls(y, z, w, u) \wedge v \leq w \Rightarrow sls(x, z, l, u)$$

Flows

Key idea: encode global data invariants as local conditions on the *flow* of nodes, an inductively computed quantity.

Example specification: nodes reachable from **root** form a tree
Solution: compute number of paths from **root** to each node

Start with a *flow domain* $(D, \Xi, +, \cdot, 0, 1)$ – here use \mathbb{N} .

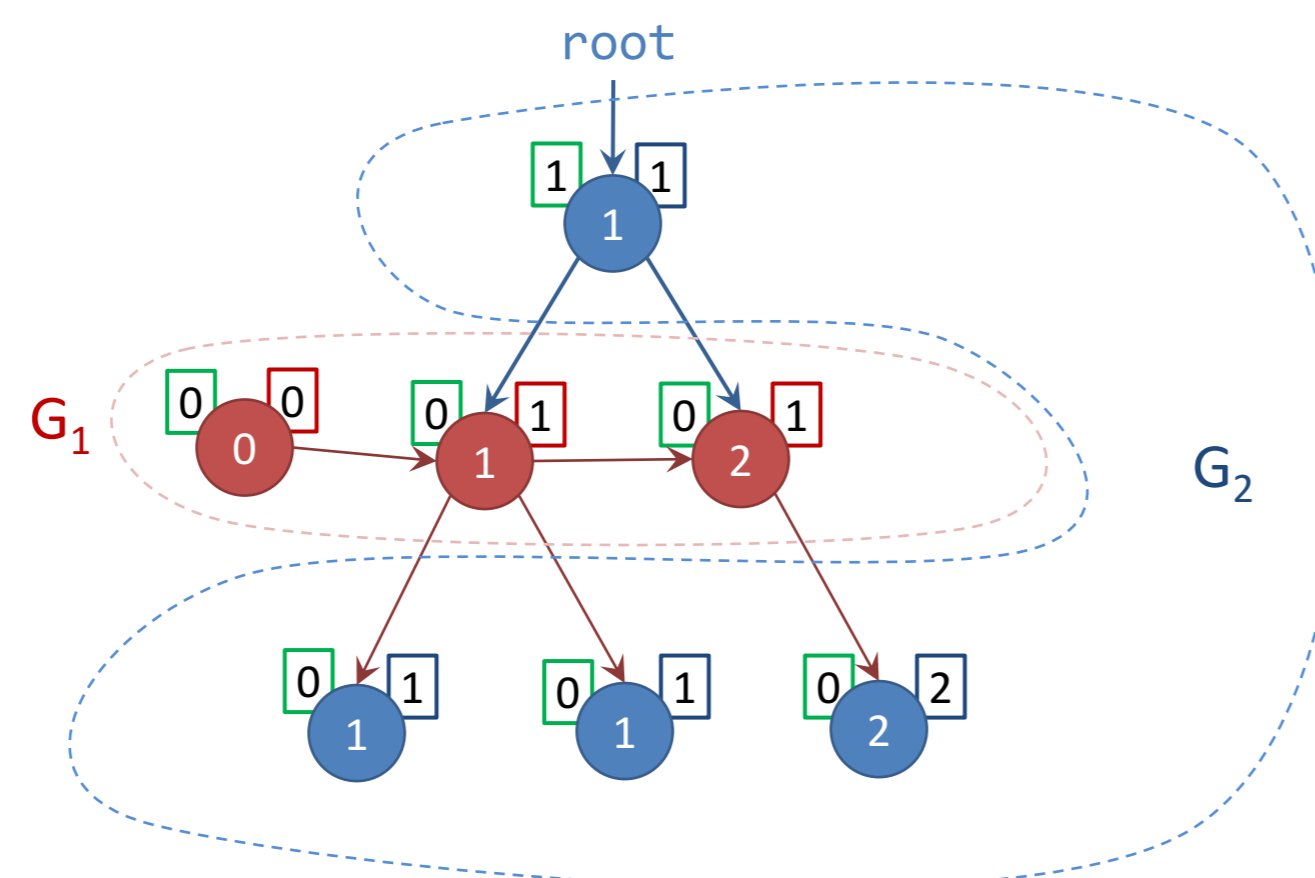
$G = (N, e)$ is a *flow graph*

- N : finite set of nodes
- e : labels edges from D

Given an inflow $in: N \rightarrow D$, compute $flow(in, G) : N \rightarrow D$

$$flow(in, G) = \text{lfp} (\lambda C. \lambda n \in N. in(n) + \sum_{n' \in N} C(n') \cdot e(n', n))$$

Example spec is now: $\forall n \in N. flow(in, G)(n) \leq 1$




Abstractions: Flow Interfaces

- Flow map of a flow graph:

$$f = fm(G)(n, n_o) = \sum_{p: n \rightsquigarrow n_o} \text{pathproduct}(p)$$

- $I = (in, f)$ is a *flow interface*
- Lift composition to interfaces: $I_1 \oplus I_2$

- $\llbracket (in, f) \rrbracket_{\text{good}}$ denotes all (in, G) s.t.
 - f is flow map of G
 - $\forall n \in G. \text{good}(n, flow(in, G)(n), G|_n)$ holds
- Example:
 - $\text{good}(n, p, _) = p \leq 1$



Flow Interface Algebras


(in, G) is a *flow interface graph*

- G : partial flow graph with outgoing edges
- in : inflow on G

Composition and decomposition:

- Defined inductively to preserve flows
- Example: $(in, G) = (in_1, G_1) \circ (in_2, G_2)$

(Flow interface graphs, \circ) is a **separation algebra**
 \Rightarrow Can use as semantic model for SL



Some nice properties:

- \oplus is associative & commutative
- $\llbracket I_1 \rrbracket \circ \llbracket I_2 \rrbracket \subseteq \llbracket I_1 \circ I_2 \rrbracket$

Highlights

- Separation logic based abstraction
- Handles unbounded sharing & overlays
- Local reasoning for shape and data
- Not tied to one traversal strategy
- Data-structure-agnostic composition and abstraction lemmas
- Simple correctness proofs for complex concurrent dictionary algorithms

Logic & Entailments

- Can use any concurrent SL-like logic
- To demonstrate, we use rely-guarantee separation logic (RGSep)
- We add new predicates
 - These are parametrized by the **good** condition

$Gr(I)$ Graph region satisfying interface I

$N(n, I)$ Singleton graph at n satisfying I

- Generic composition and decomposition:

$$\frac{Gr(I) \wedge x \in I^{in}}{N(x, I_1) * Gr(I_2) \wedge I \in I_1 \oplus I_2} \quad (\text{DECOMP})$$

$$\frac{Gr(I_1) * Gr(I_2) \wedge I \in I_1 \oplus I_2}{Gr(I)} \quad (\text{COMP})$$

Application: Verifying Concurrent Dictionaries

We can prove memory safety and linearizability of

- Harris' non-blocking singly linked list
 - B+ trees with give-up based fine grained locking
- Both use same flow abstraction and key invariants for linearizability

Example: spec of B+ tree split method:

$$\left\{ \begin{array}{l} (N(p, I_p) * N(c, I_c) \text{ --* } Gr(I)) \wedge I^{in} = \{r \mapsto (KS, 1)\} \cdot 0 \\ \wedge I^f = \epsilon \wedge I_p^f(p, c) \neq (\emptyset, 0) \wedge I_p^\alpha = (C_p, t) \wedge I_c^\alpha = (C_c, t) \end{array} \right\}$$

split(c, p);

$$\left\{ \begin{array}{l} (N(p, I_p') * N(c, I_c') * N(n, I_n)) \text{ --* } Gr(I) \wedge I^{in} = \{r \mapsto (KS, 1)\} \cdot 0 \\ \wedge I^f = \epsilon \wedge I_p'^\alpha = (C_p, t) \wedge I_c'^\alpha = (C_c', t) \wedge I_n'^\alpha = (C_n, t) \wedge C_c = C_c' \cup C_n \end{array} \right\}$$