

Flow Interfaces

Compositional Abstractions of Concurrent Data Structures

Siddharth Krishna,
Dennis Shasha, and Thomas Wies


NYU | COURANT 

Background


Verifying programs, separation logic, inductive predicates

Verifying Data Structures

```
procedure delete(x: Node)
{
  if (x != null) {
    var y := x.next;
    delete(y);
    free(x);
  }
}
```



Proof by Hand-Waving



```
procedure delete(x: Node)
```

```
{
```

```
  if (x != null) {
```

```
    var y := x.next;
```


```
    delete(y);
```

```
    free(x);
```


```
  }
```

```
}
```

```
{ }
```




Proof by Hand-Waving



Separation Logic by Example

- Points-to predicate

$$x \mapsto y$$



Stack

x	10
y	42
...	

Heap


10	42
...	
42	?

A partial heap consisting of one allocated cell

Separation Logic by Example

- Points-to predicate

$$x \mapsto y$$




Points-to predicate expresses permission to access (i.e. read/write/deallocate) heap location x **and nothing else!**

SL assertions describe the part of the heap that a program is allowed to work with.

Separation Logic by Example

- Points-to predicate

$$y \mapsto x$$



Stack

x	10
y	42
...	

Heap


10	?
...	
42	10

A partial heap consisting of one allocated cell

Separation Logic by Example

- Separating conjunction

$$x \mapsto y * y \mapsto x$$



Stack

x	10
y	42
...	

Heap


10	42
...	
42	10

Composition of
disjoint partial heaps

Separation Logic by Example

- Equalities

$$x \mapsto y \wedge x = z$$



Stack

x	10
y	42
z	10

Heap

10	42
...	
42	?

Equalities only
constrain the stack

Separation Logic by Example

- Separating conjunction

$$x \mapsto y * x \mapsto z$$


unsatisfiable?

Subheaps must be disjoint
(x can't be at two different places at once)

Separation Logic by Example

- Classical conjunction

$$x \mapsto y \wedge y \mapsto x$$



Separation Logic by Example

- Separating conjunction

$$x \mapsto z_1 * y \mapsto z_2 \wedge x = y$$

still unsatisfiable?

Convention: \wedge has higher precedence than $*$

Inductive Predicates

$$ls(x) \stackrel{\text{def}}{=} x = \text{null} \wedge \text{emp} \vee \exists y. x \mapsto y * ls(y)$$

$$ls(x) \rightsquigarrow \exists y. \quad x \mapsto y * ls(y)$$

$$\rightsquigarrow \exists y, z. \quad x \mapsto y * y \mapsto z * ls(z)$$

$$\rightsquigarrow \exists y, z, w. \quad x \mapsto y * y \mapsto z * z \mapsto w * ls(w)$$

$$\rightsquigarrow \exists y, z, w. \quad x \mapsto y * y \mapsto z * z \mapsto w * w = \text{null} \wedge \text{emp}$$

$$\rightsquigarrow \exists y, z, w. \quad x \mapsto y * y \mapsto z * z \mapsto \text{null}$$



Proof by SL

$\{ls(x)\}$

procedure delete(x: Node)

{

if (x != null) {

var y := x.next;

 delete(y);

free(x);

 }

}

$\{emp\}$

$\{ls(x) \wedge x \neq null\}$

$\{x \mapsto y * ls(y)\}$

$\{x \mapsto y * emp\}$


$\{emp\}$

Background


The Trouble with Inductive Predicates

Concurrent data structures are complex

Harris' Non-blocking List




Harris' Non-blocking List



Limitations of Inductive Predicates

Traversals need to visit each node exactly once.

$$ls(mh, null) * ls(fh, null)$$




Overlays



Limitations of Inductive Predicates

Traversals need to visit each node exactly once.

$harris(mh, fh, null) \stackrel{\text{def}}{=} \dots * harris(\dots)$




Unbounded sharing



Limitations of Inductive Predicates


$harris(mh, fh, null) \stackrel{\text{def}}{=} ls(mh, null) * \dots$



Threads can enter main list at arbitrary points


Other Approaches

Overlapping conjunction: ✱




Potentially cyclic \Rightarrow complex predicate

Updates: ramifications, reason about \rightarrow ✱



Other Approaches

Iterated separating conjunction: $\bigotimes_{x \in X} \phi(x)$



Can't do better than closed set of nodes
Every node reachable – memory leaks



Other Issues

Entailment lemmas are specific to data structure.

Example:

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

but

$$sls(x, y) * sls(y, z) \not\Rightarrow sls(x, z)$$

Solution: add lower and upper bounds $sls(x, y, l, u)$

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

Different composition lemma!

The Problem

An abstraction mechanism that can handle data structures like the Harris list?

(i.e. handle overlays, sharing, arbitrary traversals & have easy reasoning)

Background

The Trouble with Inductive Predicates

Flow Interfaces

A new abstraction for concurrent data structures


The Idea

- Inductive predicates:
 - Pro: inductive properties
 - Con: fixed traversals
- Iterated *: $\bigotimes_{x \in X} \phi(x)$
 - Pro: easy reasoning
 - Con: only local properties
- Best of both?
- Inductive properties \rightarrow local conditions
- But allow dependence on inductive quantity: **flow**

The Idea

- Express all invariants w.r.t. local condition on nodes and **flow**.
- Introduce a graph composition that preserves local invariants on global flows.
- Introduce a generic predicate to abstract a heap region satisfying the local flow condition (the **flow interface**).

Local Data Structure Invariants with Flows




Path counting!

Can we express the property that **root** points to a tree as a local condition of each node in the graph?

Flows

Step 2: Calculate the flow



Given an inflow $in: N \rightarrow D$


Here

$$in(n) = ITE(n = \text{root}, 1, 0)$$

Flows

Step 2: Calculate the flow

Start with
 $\text{flow}(n) = \text{in}(n)$




And iterate until fixpoint

$$\text{flow}(n) = \text{in}(n) + \sum_{n'} \text{flow}(n') \cdot e(n', n)$$

Flows


Step 2: Calculate the flow

Different inflows result in different flows



Flows

Step 3: Define invariant on flow



If every node satisfies the **good condition**

$$\gamma_t \stackrel{\text{def}}{=} \text{flow}(in, G)(n) = 1$$


for $in(n) = \text{ITE}(n = \text{root}, 1, 0)$

then G is “a tree rooted at **root**”

Flows

Step 3: Define invariant on flow

root



For lists, enforce at most 1 outgoing edge

$$\gamma_l \stackrel{\text{def}}{=} \text{flow}(in, G)(n) = 1 \\ \wedge (e_n = \{(n, n') \mapsto 1\} \vee e_n = \epsilon)$$

e_n : edge function restricted to nonzero edges leaving n

Flows

Alternate view

Flow graph $G = (N, E)$

Viewing E as a matrix,

$$E^2_{nn'} = \sum_m E_{nm} E_{mn'} = \text{number of 2-length paths from } n \text{ to } n'$$

$$E^l_{nn'} = \text{number of } l\text{-length paths from } n \text{ to } n'$$


$$\text{Capacity } C = I + E + E^2 + \dots$$

(converges if D is a **flow domain**)

$$\text{cap}(G)(n, n') \stackrel{\text{def}}{=} C_{nn'} = \text{the number of paths from } n \text{ to } n'$$

$$\text{flow}(in, G)(n) \stackrel{\text{def}}{=} \sum_{n'} in(n') \cdot \text{cap}(G)(n', n)$$

Data Invariants




- Flow domain: ω -cpo & positive semiring
($D, \sqsubseteq, \sqcup, +, \cdot, 0, 1$)
- Another Example: lower-bound domain
($\mathbb{Z} \cup \{-\infty, \infty\}, \geq, \min, \min, \max, \infty, -\infty$)
- Label each edge with value of source node
- Use $in(n) = ITE(n = \text{root}, -\infty, \infty)$
- $\text{flow}(in, G)(n) = \max$ value in *some* path from **root** to n
- These paths are sorted if

$$\gamma_S \stackrel{\text{def}}{=} \text{flow}(in, G)(n) \leq a$$

where a is value at n

Stacking Flows




- Want shape & data invariants?
- Flow of product is product of flows!
- Example: min-heap
- Use product of path-counting and lower-bound domain
- Use $in(n) = ITE(n = \text{root}, (1, -\infty), (0, \infty))$
- And good condition

$$\gamma_n \stackrel{\text{def}}{=} \text{flow}(in, G)(n) = (1, l) \wedge l \leq a$$

where a is value at n

Stacking Flows



- Data invariants are decoupled from shape invariants

- Min-heap


$$\gamma_h \stackrel{\text{def}}{=} \text{flow}(in, G)(n) = (1, l) \wedge l \leq a$$

- Sorted list

$$\gamma_{sl} \stackrel{\text{def}}{=} \text{flow}(in, G)(n) = (1, l) \wedge l \leq a \\ \wedge (e_n = \{(n, n') \mapsto 1\} \vee e_n = \epsilon)$$

Harris List

- We can now describe Harris' List
- Flow domain: two path-counting flows
 - One from mh and one from fh
 - Every node is on at least one of these lists
- Nodes labelled: marked/unmarked
 - All nodes in free list are marked



Expressivity of Flows


- Flows can describe:
 - Lists (singly and doubly linked, cyclic)
 - Trees (n-ary, arbitrary arity)
 - Nested combinations
 - Sorted lists, binary heaps, BSTs
 - Overlaid structures (threaded and B-link trees)
 - Irregular structures (DAGs, graphs)
 - Unbounded sharing & irregular traversals (Harris)
- But inductive predicates easier for:
 - Simple inductive structures/abstractions

Compositional Reasoning

But can we reason compositionally about flows and graphs à la SL?

Graph Composition

- Standard SL Composition (disjoint union) is too weak:




- Need a composition that preserves flow!

Flow Graphs

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

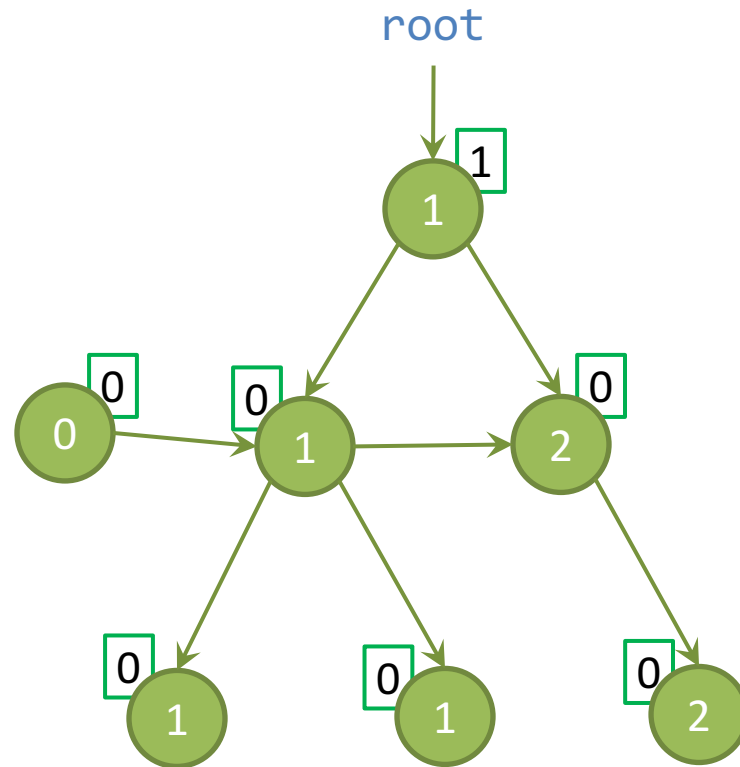
- $G = (N, N_o, e)$ is a partial graph with
 - N the set of internal nodes of the graph
 - N_o the set of *sink* nodes of the graph
 - $e: N \times (N \cup N_o) \rightarrow D$ is an edge function
- $in: N \rightarrow D$ is an inflow



Inflow in specifies **rely** of G on its context.


Flow Graph Composition

(in, G)



Flow Graph Composition

$$(in, G) = (in_1, G_1) \circ (in_2, G_2)$$



$$\text{If } in_1(n) = in(n) + \sum_{n' \in G_2} \text{flow}(in, G)(n') \cdot e(n', n)$$

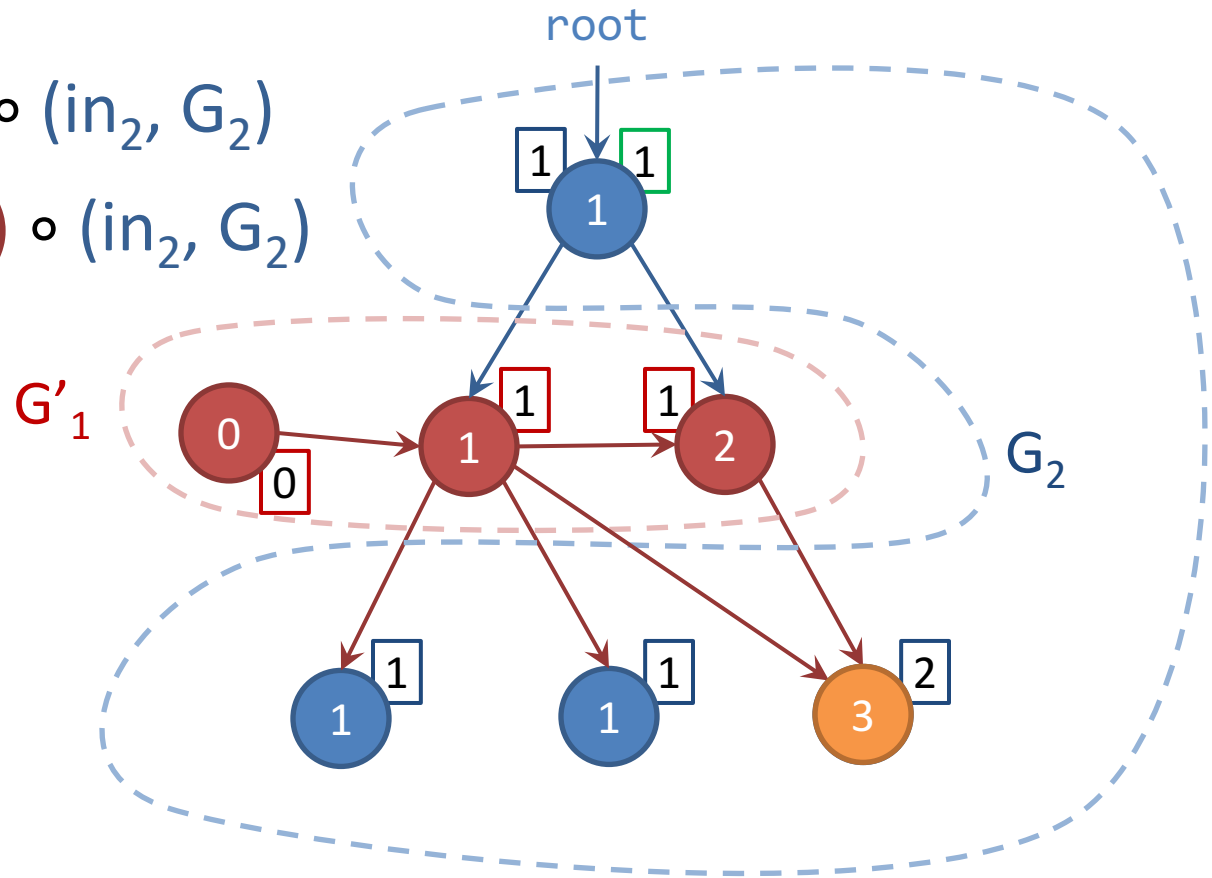
Flow Graph Composition

- $H_1 \circ H_2$ is
 - commutative: $H_1 \circ H_2 = H_2 \circ H_1$
 - associative : $(H_1 \circ H_2) \circ H_3 = H_1 \circ (H_2 \circ H_3)$
 - cancelative: $H \circ H_1 = H \circ H_2 \Rightarrow H_1 = H_2$
- ⇒ Flow graphs form a [separation algebra](#).
- ⇒ We can use them to give semantics to SL assertions.
- How do we abstract flow graphs?

Modifying the Graph


$$(in, G) = (in_1, G_1) \circ (in_2, G_2)$$

$$(in, G') \neq (in_1, G'_1) \circ (in_2, G_2)$$



Need to preserve the flow into G_2 ...


Flow Map of a Flow Graph




$$\text{flm}(in, G)(n, n_o) = \sum e(n, n_1) \cdots e(n_k, n_o) \quad \text{over all paths in } G$$

$$\text{flow}(in, G)(n_o) = \sum_{n \in G} in(n) \cdot \text{flm}(in, G)(n, n_o)$$

Flow Map: Example




Flow Map: Example



Flow map abstracts from internal structure of the graph


Flow Map: Example



Flow map abstracts from internal structure of the graph

Flow Interfaces

- $I = (in, a, f)$ is a **flow interface**
 - $in: N \rightarrow D$ is an inflow
 - $a \in A$ is the abstraction of node labels
 - $f: N \times N_o \rightarrow D$ is a flow map



$$\in \llbracket (\{r \rightsquigarrow 1, \dots\}, _, \{(r, y) \rightsquigarrow 1\}) \rrbracket$$

Flow Interfaces

- Flow interfaces induce a congruence on flow graphs:

$$H_i, H'_i \in \llbracket I_i \rrbracket \wedge H_1 \circ H_2 \in \llbracket I \rrbracket \Rightarrow H'_1 \circ H'_2 \in \llbracket I \rrbracket$$

- Lift flow graph composition to interfaces: \oplus
- (Flow Interfaces, \oplus) also a separation algebra!

Reasoning about Modifications

- \oplus congruence \Rightarrow can replace G_1 with any G'_1 with same interface




- Showing equivalence: requires fixpoint reasoning
- But concurrent algos modify small regions

Separation Logic with Flow Interfaces

- Good graph predicate $Gr_\gamma(\mathbf{I})$
 - γ : SL predicate that defines good node condition and abstraction of heap onto nodes of flow graph
 - \mathbf{I} : flow interface
- Good node predicate $N_\gamma(x, \mathbf{I})$
 - like Gr but denotes a single node
 - definable in terms of Gr
- Dirty region predicate $[P]_{\gamma, \mathbf{I}}$
 - P : SL predicate
 - denotes heap region that is expected to satisfy interface \mathbf{I} but may currently not

Example

Sorted Linked List



$$\begin{aligned} \gamma(x, in, a, f) &\stackrel{\text{def}}{=} n \mapsto k, n' \wedge in(n) = (1, l) \\ &\wedge a = \{k\} \wedge l \leq k \\ &\wedge f = ITE(n' = null, \epsilon, \{(n, n') \mapsto 1\}) \end{aligned}$$

Global invariant:

$$Gr(I) \wedge I^{in} + \mathbf{0} = \{r \mapsto (1, -\infty)\} + \mathbf{0} \wedge I^f = \epsilon$$

Harris' List

$$\begin{aligned} \gamma(n, in, a, f) := & \exists n', n''. n \mapsto n', n'' \wedge a \neq \top \wedge (M(n') \Leftrightarrow a \neq \diamond) \wedge (0, 0) < in(n) \leq (1, 1) \\ & \wedge (in(n) \geq (0, 1) \Rightarrow a \neq \diamond) \wedge (n = ft \Rightarrow in(n) \geq (0, 1)) \wedge (in(n) \leq (1, 0) \Rightarrow n'' = null) \\ & \wedge f = \text{ITE}(u(n') = null, \epsilon, \{(n, u(n')) \succrightarrow (1, 0)\}) + \text{ITE}(n'' = null, \epsilon, \{(n, n'') \succrightarrow (0, 1)\}). \end{aligned}$$

Global invariant:

$$\Phi \stackrel{\text{def}}{=} \exists I. \text{Gr}(I) \wedge I^{in} + \mathbf{0} = \{mh \succrightarrow (1, 0)\} + \{fh \succrightarrow (1, 0)\} + \mathbf{0} \wedge I^f = \epsilon$$

Data-Structure-Agnostic Proof Rules

- Decomposition

$$\text{Gr}(I) \wedge x \in I \models \exists I_1, I_2. \text{N}(x, I_1) * \text{Gr}(I_2) \wedge I = I_1 \oplus I_2$$

- Step

$$I = I_1 \oplus I_2 \wedge (x, y) \in I_1^f \wedge I^f = \epsilon \models y \in I_2$$

- Composition

$$\text{Gr}(I_1) * \text{Gr}(I_2) \models \text{Gr}(I_1 \oplus I_2)$$

Harris: Insert

```
procedure insert() {  $\{\Phi\}$   
  var l := mh;  $\longleftrightarrow \{\text{Gr}(I) \wedge mh \in I^{in}\}$   
  var r := unmarked(l.next);  
  while (r != null && ?? ) {  $\{\text{Gr}(I) \wedge mh \in I^{in}\} \xrightarrow{\text{(Decomp)}} \{N(l, I_l) * \text{Gr}(I_2) \wedge I = I_l \oplus I_2\}$   
    l := r;  
    r := unmarked(l.next);  
  }  
  ...  
}
```

Harris: Insert


```
procedure insert() {  $\{\Phi\}$   
  var l := mh;  $\{N(l, I_l) * \text{Gr}(I_2) \wedge I = I_l \oplus I_2\}$   
  var r := unmarked(l.next);  
  while (r != null && ?? ) {  $\{I = I_l \oplus I_2 \wedge (l, r) \in I_l^f \wedge I^f = \epsilon\}$   
    l := r;  
    r := unmarked(l.next);  
  }  
  ...  
}
```

\downarrow (Step), (Decomp)

$\{N(l, I_l) * N(r, I_r) * \text{Gr}(I_3) \wedge I = I_l \oplus I_r \oplus I_3\}$

Harris: Insert

```
procedure insert() {  $\{\Phi\}$   
  var l := mh;  $\{N(l, I_l) * Gr(I_2) \wedge I = I_l \oplus I_2\}$   
  var r := unmarked(l.next);  
  while (r != null && ?? ) {  $\left\{ \begin{array}{l} N(l, I_l) * N(r, I_r) * Gr(I_3) \\ \wedge I = I_l \oplus I_r \oplus I_3 \end{array} \right\}$   
    l := r;  
    r := unmarked(l.next);  
  }  
  ...  
}
```

 (Comp)

$\{N(r, I_r) * Gr(I_4) \wedge I = I_r \oplus I_4\}$

Harris: Insert

```
procedure insert() {  $\{\Phi\}$ 
  var l := mh;  $\{N(l, I_l) * \text{Gr}(I_2) \wedge I = I_l \oplus I_2\}$ 
  var r := unmarked(l.next);
  while (r != null && ?? ) {
    l := r;  $\{N(l, I_l) * \text{Gr}(I_4) \wedge I = I_l \oplus I_4\}$ 
    r := unmarked(l.next);
  }
  ...
}
```

Flow Interfaces

- Data structure abstractions that
 - can handle unbounded sharing and overlays
 - treat structural and data constraints uniformly
 - do not encode specific traversal strategies
 - provide data-structure-agnostic composition and decomposition rules
 - remain within general theory of separation logic

Background

The Trouble with Inductive Predicates

Flow Interfaces


Concurrent Dictionaries

A memory-safe, linearizable, abstract template

Concurrent Dictionaries

- Dictionary: key-value store
- Examples: sorted linked lists, BSTs, B-trees, hash maps, ...
- With flows: generic template + proof

Inset Flows




KS: the set of all search keys
e.g. $KS = \text{Int}$

Inset flow domain:
 $(2^{KS}, \subseteq, \cup, \cap, \emptyset, KS)$

Label each edge with the set of keys that follow that edge in a search (edgeset).

Inset Flows




KS: the set of all search keys
e.g. $KS = \text{Int}$

Inset flow domain:
 $(2^{KS}, \subseteq, \cup, \cap, \emptyset, KS)$

Set inflow *in* of **root** to KS and to \emptyset for all other nodes.

Inset Flows



$$I_1 = \{v \mid 3 < v\}$$

$$I_2 = \{v \mid 3 < v < 6\}$$


$$I_3 = \{v \mid 8 < v\}$$

$\text{flow}(\text{in}, G)(n)$ is the *inset* of node n , i.e., the set of keys k such that a search for k will traverse node n .

From Insets to Keysets

$$\text{outset}(G)(n) =$$

$$\bigcup_{n \in N} e(n, n')$$



$$\text{keyset}(\text{in}, G)(n) =$$

$$\text{inset}(\text{in}, G)(n) \setminus \text{outset}(G)(n)$$

$\text{keyset}(\text{in}, G)(n)$ is the set of keys that could be in n

Verifying Concurrent Dictionaries

- **Good state** conditions
 - edgesets are disjoint for each n :
 $\{e(n, n')\}_{n' \in N}$ are disjoint
 - keyset of each n covers n 's contents:
 $C(G)(n) \subseteq \text{keyset}(\text{in}, G)(n)$
- **Keyset Theorem** [Shasha and Goodman, 1988]

If all ops preserve good state, and methods
search/insert/delete k at node s.t. $k \in \text{keyset}(n)$

\Rightarrow The implementation is linearizable

Encoding Using Flows

Description of good node

Description of locked node

$$\gamma(x, in, (C, T), f) := \exists t. (\underbrace{\gamma_g(x, in, C, t, f) \wedge T = \{t\}}_{\text{Contents are in keyset}} \vee \underbrace{\gamma_b(x, t) \wedge t \neq 0 \wedge T = \{\bar{t}\}}_{\text{Edgesets leaving node are disjoint}}) \\ \wedge C \subseteq in(x) \wedge \forall y. (C \cap f(x, y) = \emptyset \wedge \forall z. f(x, y) \cap f(x, z) = \emptyset)$$

Contents are in keyset

Edgesets leaving node are disjoint

Global invariant:

$$\Phi := \exists I, in. Gr(I) \wedge in \in I^{In} \wedge in + \mathbf{0} = \{r \mapsto KS\} + \mathbf{0} \wedge I^f = \epsilon$$

Give-up Template

```
1 procedure dictionaryOp(Key k) {  $\{\Phi\}$ 
2   var c := r;  $\{N(c, I_c) \dashv\!\!\dashv \Phi\}$ 
3   while (true) {  $\{N(c, I_c) \dashv\!\!\dashv \Phi\}$ 
4     lock(c);  $\{N(c, I_c) \dashv\!\!\dashv \Phi \wedge I_c^a = (\_, \{t\})\}$ 
5     var n;
6     if (inRange(c, k)) {  $\{N(c, I_c) \dashv\!\!\dashv \Phi \wedge I_c^a = (\_, \{t\}) \wedge k \in I_c^{In}(c)\}$ 
7       n := findNext(c, k);  $\left\{ \begin{array}{l} N(c, I_c) \dashv\!\!\dashv \Phi \wedge I_c^a = (\_, \{t\}) \wedge k \in I_c^{In}(c) \\ \wedge (n \neq \text{null} \wedge k \in I_c^f(c, n) \vee n = \text{null} \wedge \forall x \in I_c^f. k \notin I_c^f(c, x)) \end{array} \right\}$ 
8     if (n == null) break;
9      $\{(N(c, I_c) * N(n, I_n)) \dashv\!\!\dashv \Phi \wedge I_c^a = (\_, \{t\})\}$ 
10  } else {
11    n := r;  $\{N(c, I_c) \dashv\!\!\dashv \Phi \wedge I_c^a = (\_, \{t\}) \wedge n = r\}$ 
12  }  $\{(N(c, I_c) * N(n, I_n)) \dashv\!\!\dashv \Phi \wedge I_c^a = (\_, \{t\}) \vee N(c, I_c) \dashv\!\!\dashv \Phi \wedge I_c^a = (\_, \{t\}) \wedge c = n = r\}$ 
13  unlock(c);
14  c := n;  $\{N(c, I_c) \dashv\!\!\dashv \Phi\}$ 
15 }  $\{N(c, I_c) \dashv\!\!\dashv \Phi \wedge I_c^a = (\_, \{t\}) \wedge k \in I_c^{In}(c) \wedge \forall x \in I_c^f. k \notin I_c^f(c, x)\}$ 
16 var res := decisiveOp(c, k);  $\{N(c, I_c) \dashv\!\!\dashv \Phi \wedge I_c^a = (\_, \{t\})\}$ 
17 unlock(c);  $\{N(c, I_c) \dashv\!\!\dashv \Phi\}$ 
18 return res;  $\{\Phi\}$ 
19 }
```

Actions

$$t \in T \wedge N(x, (In, (C, \{0\}), f)) \rightsquigarrow N(x, (In, (C, T'), f)) \wedge T' \subseteq \{t, \bar{t}\} \quad (\text{Lock})$$

$$t \in T \wedge emp \rightsquigarrow N(x, (\{\{x \mapsto 0\}\}, (\emptyset, \{\bar{t}\}), \epsilon)) \quad (\text{Alloc})$$

$$t \in T \wedge Gr(I) \wedge I^a \sqsubseteq (_, \{t, \bar{t}\}) \rightsquigarrow Gr(I') \wedge I'^a \sqsubseteq (_, \{0, t, \bar{t}\}) \wedge I \preceq I' \quad (\text{Sync})$$

Keyset Theorem with Flows

Theorem: An implementation of the template with appropriate γ_g, γ_b that satisfy the spec below (with some side conditions) is memory safe and linearizable.

$$\left\{ \begin{array}{l} \boxed{N(c, I_c) \dashv\ast \Phi} \wedge I_c^a = (C, \{t\}) \wedge k \in I_c^{In}(c) \\ \wedge \forall x \in I_c^f. k \notin I_c^f(c, x) \end{array} \right\} \text{res} := \text{decisiveOp}(c, k); \left\{ \boxed{N(c, I'_c) \dashv\ast \Phi} \wedge I_c \approx I'_c \wedge \Psi \right\}$$

$$\text{where } \Psi := \begin{cases} I'_c{}^a = (C, \{t\}) \wedge \text{res} \Leftrightarrow k \in C & \text{for member} \\ I'_c{}^a = (C \cup \{k\}, \{t\}) \wedge \text{res} \Leftrightarrow k \notin C & \text{for insert} \\ I'_c{}^a = (C \setminus \{k\}, \{t\}) \wedge \text{res} \Leftrightarrow k \in C & \text{for delete} \end{cases}$$

Background

The Trouble with Inductive Predicates

Flow Interfaces

Concurrent Dictionaries

Conclusion

A new way to reason about data structures

Conclusion

- Radically new approach for building compositional abstractions of data structures.
- Fits in existing (concurrent) separation logics.
- Enables simple correctness proofs of concurrent data structure algorithms.
- Proofs can abstract from the details of the specific data structure implementation.