

Local Reasoning for Global Graph Properties^{*}

Siddharth Krishna¹, Alexander J. Summers², and Thomas Wies¹

¹ New York University, New York, NY, USA, {siddharth,wies}@cs.nyu.edu

² ETH Zürich, Zurich, Switzerland, alexander.summers@inf.ethz.ch

Abstract. Separation logics are widely used for verifying programs that manipulate complex heap-based data structures. These logics build on so-called *separation algebras*, which allow expressing properties of heap regions such that modifications to a region do not invalidate properties stated about the remainder of the heap. This concept is key to enabling modular reasoning and also extends to concurrency. While heaps are naturally related to mathematical graphs, many ubiquitous graph properties are non-local in character, such as reachability between nodes, path lengths, acyclicity and other structural invariants, as well as data invariants which combine with these notions. Reasoning modularly about such graph properties remains notoriously difficult, since a local modification can have side-effects on a global property that cannot be easily confined to a small region.

In this paper, we address the question: What separation algebra can be used to avoid proof arguments reverting back to tedious global reasoning in such cases? To this end, we consider a general class of global graph properties expressed as fixpoints of algebraic equations over graphs. We present mathematical foundations for reasoning about this class of properties, imposing minimal requirements on the underlying theory that allow us to define a suitable separation algebra. Building on this theory, we develop a general proof technique for modular reasoning about global graph properties expressed over program heaps, in a way which can be directly integrated with existing separation logics. To demonstrate our approach, we present local proofs for two challenging examples: a priority inheritance protocol and the non-blocking concurrent Harris list.

1 Introduction

Separation logic (SL) [34,41] provides the basis of many successful verification tools that can verify programs manipulating complex data structures [1, 5, 20, 31]. This success is due to the logic’s support for reasoning modularly about modifications to heap-based data. For simple inductive data structures such as lists and trees, much of this reasoning can be automated [2, 12, 23, 36]. However, these techniques often fail when data structures are less regular (e.g. multiple overlaid data structures) or provide multiple traversal patterns (e.g. threaded trees). Such idioms are prevalent in real-world implementations such as the fine-grained concurrent data structures found in operating systems and databases. Solutions to these problems have been proposed [16] but remain difficult to automate. For proofs of general graph algorithms, the situation is even more dire. Despite substantial improvements in the verification methodology for such algorithms [39,42],

^{*} This is an extended version of a paper that appeared at ESOP 2020.

```

1 method acquire(p: Node, r: Node) {
2   if (r.next == null) {
3     r.next := p; update(p, -1, r.curr_prio)
4   } else {
5     p.next := r; update(r, -1, p.curr_prio)
6   }
7 }
8 method update(n: Node, from: Int, to: Int) {
9   n.prios := n.prios \ {from}
10  if (to >= 0) n.prios := n.prios ∪ {to}
11  from := n.curr_prio
12  n.curr_prio := max(n.prios ∪ {n.def_prio})
13  to := n.curr_prio;
14  if (from != to && n.next != null) {
15    update(n.next, from, to)
16  }
17 }

```

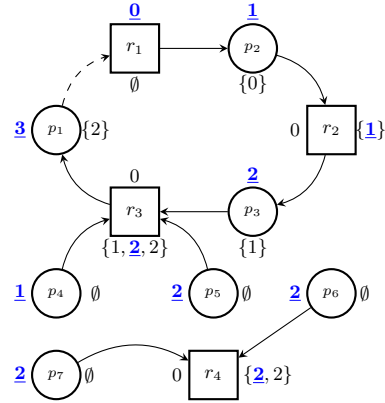


Fig. 1: Pseudocode of the PIP and a state of the protocol data structure. Round nodes represent processes and rectangular nodes resources. Nodes are marked with their default priorities `def_prio` as well as the aggregate priority multiset `prios`. A node’s current priority `curr_prio` is underlined and marked in bold blue.

significant parts of the proof argument still typically need to be carried out using non-local reasoning [8, 9, 14, 27]. This paper presents a general technique for local reasoning about global graph properties that can be used within off-the-shelf separation logics. We demonstrate our technique using two challenging examples for which no fully local proof existed before, respectively, whose proof required a tailor-made logic.

As a motivating example, we consider an idealized priority inheritance protocol (PIP), a technique used in process scheduling [43]. The purpose of the protocol is to avoid *priority inversion*, i.e. a situation where a low-priority process causes a high-priority process to be blocked. The protocol maintains a bipartite graph with nodes representing processes and resources. An example graph is shown in Fig. 1. An edge from a process p to a resource r indicates that p is waiting for r to be available whereas an edge in the other direction means that r is currently held by p . Every node has an associated *default* priority and *current*; these are natural numbers. The current priority is used for scheduling processes. When a process attempts to acquire a resource currently held by another process, the graph is updated to avoid priority inversion. For example, when process p_1 with current priority 3 attempts to acquire the resource r_1 held by process p_2 of priority 1, p_1 ’s higher priority is propagated to p_2 and, transitively, to any other process that p_2 is waiting for (p_3 in this case). As a result, all nodes on the created cycle³ will get current priority 3. The protocol maintains the following *invariant*: the current priority of each node is the maximum of its default priority and the current priorities of all its predecessors. Priority propagation is implemented by the method `update` shown in Fig 1. The implementation represents graph edges by `next` pointers and handles both adding an edge (`acquire`) and removing one (`release` - code omitted). To recalculate the current priority of a node (line 12), each node maintains its default priority `def_prio` and a multiset `prios` which contains the priorities of all its immediate predecessors.

³ The cycle can be used to detect/handle a deadlock; this is not the concern of this data structure.

Verifying that the PIP maintains its invariant using established separation logic (SL) techniques is challenging. In general, SL assertions describe resources and express the fact that the program has permission to access and manipulate these resources. In what follows, we stick to the standard model of SL where resources are memory regions represented as partial heaps. We sometimes view partial heaps more abstractly as partial graphs (hereafter, simply graphs). Assertions describing larger regions are built from smaller ones using *separating conjunction*, $\phi_1 * \phi_2$. Semantically, the $*$ operator is tied to a notion of resource composition defined by an underlying *separation algebra* [6, 7]. In the standard model, composition enforces that ϕ_1 and ϕ_2 must describe disjoint regions. The logic and algebra are set up so that changes to the region ϕ_1 do not affect ϕ_2 (and vice versa). That is, if $\phi_1 * \phi_2$ holds before the modification and ϕ_1 is changed to ϕ'_1 , then $\phi'_1 * \phi_2$ holds afterwards. This so-called *frame rule* enables modular reasoning about modifications to the heap and extends well to the concurrent setting when threads operate on disjoint portions of memory [4, 10, 11, 40]. However, the mere fact that ϕ_2 is preserved by modifications to ϕ_1 does not guarantee that if a global property such as the PIP invariant holds for $\phi_1 * \phi_2$, it also still holds for $\phi'_1 * \phi_2$.

For example, consider the PIP scenario depicted in Fig. 1. If ϕ_1 describes the subgraph containing only node p_1 , ϕ_2 the remainder of the graph, and ϕ'_1 the graph obtained from ϕ_1 by adding the edge from p_1 to r_1 , then the PIP invariant will no longer hold for the new composed graph described by $\phi'_1 * \phi_2$. On the other hand, if ϕ_1 captures p_1 and the nodes reachable from r_1 (i.e., the set of nodes modified by `update`), ϕ_2 the remainder of the graph, and we reestablish the PIP invariant locally in ϕ_1 obtaining ϕ'_1 (i.e., run `update` to completion), then $\phi'_1 * \phi_2$ will also globally satisfy the PIP invariant. The separating conjunction $*$ is not sufficient to differentiate these two cases; both describe valid partitions of a possible program heap. As a consequence, prior techniques have to revert back to non-local reasoning to prove that the invariant is maintained.

A first helpful idea towards a solution to this problem is that of *iterated separating conjunction* [32, 48], which describes a graph G consisting of a set of nodes X by a formula $\Psi = \bigstar_{x \in X} N(x)$ where $N(x)$ is some predicate that holds locally for every node $x \in X$. Using such node-local conditions one can naturally express non-inductive properties of graphs (e.g. “ G has no outgoing edges” or “ G is bipartite”). The advantages of this style of specification are two-fold. First, one can arbitrarily decompose and recompose Ψ by splitting X into disjoint subsets. For example, if X is partitioned into X_1 and X_2 , then Ψ is equivalent to $\bigstar_{x \in X_1} N(x) * \bigstar_{x \in X_2} N(x)$. Moreover, it is very easy to prove that Ψ is preserved under modifications of subgraphs. For instance, if a program modifies the subgraph induced by X_1 such that $\bigstar_{x \in X_1} N(x)$ is preserved locally, then the frame rule guarantees that Ψ will be preserved in the new larger graph. Iterated separating conjunction thus yields a simple proof technique for local reasoning about graph properties that can be described in terms of node-local conditions. However, this idea alone does not actually solve our problem because general global graph properties such as “ G is a direct acyclic graph”, “ G is an overlay of multiple trees”, or “ G satisfies the PIP invariant” cannot be directly described via node-local conditions.

Solution. The key ingredient of our approach is the concept of a *flow* of a graph: a function fl from the nodes of the graph to *flow values*. For the PIP, the flow maps each node to the multiset of its incoming priorities. In general, a flow is a fixpoint of

a set of algebraic equations induced by the graph. These equations are defined over a *flow domain*, which determines how flow values are propagated along the edges of the graph and how they are aggregated at each node. In the PIP example, an edge between nodes (n, n') propagates the multiset containing $\max(fl(n), n.\text{def_prio})$ from n to n' . The multisets arriving at n' are aggregated with multiset union to obtain $fl(n')$. Flows enable capturing global graph properties in terms of node-local conditions. For example, the PIP invariant can be expressed by the following node-local condition: $n.\text{curr_prio} = \max(fl(n), n.\text{def_prio})$. To enable compositional reasoning about such properties we need an appropriate separation algebra allowing us to prove locally that modifications to a subgraph do not affect the flow of the remainder of the graph.

To this end, we make the useful observation that a separation algebra induces a notion of an *interface of a resource*: we say that two resources a and a' are equivalent if they compose with the same resources. The interface of a resource a could then be defined as a 's equivalence class, but more-succinct and simpler representations may be possible. In the standard model of SL where resources are graphs and composition is disjoint graph union, the interface of a graph G is the set of all graphs G' that have the same domain as G ; in this model, a graph's domain could be defined to be its interface.

The interfaces of resources described by assertions capture the information that is implicitly communicated when these assertions are conjoined by separating conjunction. As we discussed earlier, in the standard model of SL, this information is too weak to enable local reasoning about global properties of the composed graphs because some additional information about the subgraphs' structure other than which nodes they contain must be communicated. For instance, if the goal is to verify the PIP invariant, the interfaces must capture information about the multisets of priorities propagated between the subgraphs. We define a separation algebra achieving exactly this: the induced *flow interface* of a graph G in this separation algebra captures how values of the flow domain must enter and leave G such that, when composed with a compatible graph G' , the imposed local conditions on the flow of each node are satisfied in the composite graph.

This is the key to enabling SL-style framing for global graph properties. Using iterated separating conjunctions over the new separation algebra, we obtain a compositional proof technique that yields succinct proofs of programs such as the PIP, whose proofs with existing techniques would involve non-trivial global reasoning steps.

Contributions. In §2, we present mathematical foundations for flow domains, imposing the minimal requirements on the underlying algebra that allow us to capture a broad range of data structure invariants and graph properties and reason locally about them in a suitable separation algebra. Building on this theory we develop a general proof technique for modular reasoning about global graph properties that can be integrated with existing separation logics (§3). We further identify general mathematical conditions that can be used when desired to guarantee unique flows, and provide local proof arguments to check the preservation of these conditions (§4). We demonstrate the versatility of our approach by presenting local proofs for two challenging examples: the PIP and the concurrent non-blocking list due to Harris [13].

Flows Redesigned. Our work is inspired by the recent flow framework explored by some of the authors [25], but was redesigned from the ground up. We revisit the core

algebra behind flow reasoning, and derive a different algebraic foundation by analysing the minimal requirements for general local reasoning; we call our newly-designed reasoning framework the *foundational flow framework*. Our new framework makes several significant improvements over [25] and eliminates its most stark limitations. We provide a detailed technical comparison with [25] and discuss other related work in §5.

2 The Foundational Flow Framework

In this section, we introduce the foundational flow framework, explaining the motivation for its design with respect to local reasoning principles. We aim for a general technique for modularly proving the preservation of recursively-defined invariants over (partial) graphs, with well-defined decomposition and composition operations.

2.1 Preliminaries and Notation

The term $(b ? t_1 : t_2)$ denotes t_1 if condition b holds and t_2 otherwise. We write $f : A \rightarrow B$ for a function from A to B , and $f : A \dashrightarrow B$ for a partial function from A to B . For a partial function f , we write $f(x) = \perp$ if f is undefined at x . We use lambda notation $(\lambda x. E)$ to denote a function that maps x to the expression E (typically containing x). If f is a function from A to B , we write $f[x \mapsto y]$ to denote the function from $A \cup \{x\}$ defined by $f[x \mapsto y](z) := (z = x ? y : f(z))$. We use $\{x_1 \mapsto y_1, \dots, x_n \mapsto y_n\}$ for pairwise different x_i to denote the function $\epsilon[x_1 \mapsto y_1] \cdots [x_n \mapsto y_n]$, where ϵ is the function on an empty domain. Given functions $f_1 : A_1 \rightarrow B$ and $f_2 : A_2 \rightarrow B$ we write $f_1 \uplus f_2$ for the function $f : A_1 \uplus A_2 \rightarrow B$ that maps $x \in A_1$ to $f_1(x)$ and $x \in A_2$ to $f_2(x)$ (if A_1 and A_2 are not disjoint sets, $f_1 \uplus f_2$ is undefined).

We write $\delta_{n=n'} : M \rightarrow M$ for the function defined by $\delta_{n=n'}(m) := m$ if $n = n'$ else 0. We also write $\lambda_0 := (\lambda m. 0)$ for the identically zero function, $\lambda_{\text{id}} := (\lambda m. m)$ for the identity function, and use $e \equiv e'$ to denote function equality. For $e : M \rightarrow M$ and $m \in M$ we write $m \triangleright e$ to denote the function application $e(m)$. We write $e \circ e'$ to denote function composition, i.e. $(e \circ e')(m) = e(e'(m))$ for $m \in M$, and use superscript notation e^p to denote the function composition of e with itself p times.

For multisets S , we use standard set notation when clear from the context. We write $S(x)$ to denote the number of occurrences of x in S . We write $\{x_1 \mapsto i_1, \dots, x_n \mapsto i_n\}$ for the multiset containing i_1 occurrences of x_1 , i_2 occurrences of x_2 , etc.

A *partial monoid* is a set M , along with a partial binary operation $+$: $M \times M \dashrightarrow M$, and a special zero element $0 \in M$, such that (1) $+$ is associative, i.e., $(m_1 + m_2) + m_3 = m_1 + (m_2 + m_3)$; and (2) 0 is an identity, i.e., $m + 0 = 0 + m = m$. Here, $=$ means either both sides are defined and equal, or both are undefined. We identify a partial monoid with its support set M . If $+$ is a total function, then we call M a monoid. Let $m_1, m_2, m_3 \in M$ be arbitrary elements of the (partial) monoid in the following. We call a (partial) monoid M *commutative* if $+$ is commutative, i.e., $m_1 + m_2 = m_2 + m_1$. Similarly, a commutative monoid M is *cancellative* if $+$ is cancellative, i.e., if $m_1 + m_2 = m_1 + m_3$ is defined, then $m_2 = m_3$.

A *separation algebra* [6] is a cancellative, partial, commutative monoid.

2.2 Flows

Recursive properties of graphs naturally depend on non-local information; e.g. we cannot express that a graph is acyclic directly as a conjunction of per-node invariants.

Our foundational flow framework defines *flow values* at each node that capture non-local graph properties, and enables local specification and reasoning about such properties. Flow values are drawn from a *flow domain*, an algebraic structure which also specifies the operations used to define a flow via recursive computations over the graph. Our entire theory is parametric with the choice of a flow domain, whose components will be explained and motivated in the rest of this section.

Definition 1 (Flow Domain). A flow domain $(M, +, 0, E)$ consists of a commutative cancellative (total) monoid $(M, +, 0)$ and a set of edge functions $E \subseteq M \rightarrow M$.

Example 1. The *path-counting* flow domain is $(\mathbb{N}, +, 0, \{\lambda_{\text{id}}, \lambda_0\})$, consisting of the monoid of natural numbers under addition and the set of edge functions containing only the identity function and the zero function. This can be used to define a flow where the values at each node represent the number of paths to this node from a distinguished node n . Path-counting provides enough information to express locally per node that e.g. (a) all nodes are reachable from n (all path counts are non-zero), or (b) that the graph forms a tree rooted at n (all path counts are exactly 1).

Example 2. We use $(\mathbb{N}^{\mathbb{N}}, \cup, \emptyset, \{\lambda_0\} \cup \{(\lambda m. \{\max(m \cup \{p\})\}) \mid p \in \mathbb{N}\})$ as flow domain for the PIP example (Figure 1). This consists of the monoid of multisets of natural numbers under multiset union and two kinds of edge functions: λ_0 and functions mapping a multiset m to the singleton multiset containing the maximum value between m and a fixed value p (used to represent a node’s default priority). This can define a flow which locally captures the appropriate current node priorities as the graph is modified.

Further definitions in this section assume a fixed flow domain $(M, +, 0, E)$ and a (potentially infinite) set of nodes \mathfrak{N} . For this section, we abstract heaps using directed partial graphs; integration of our graph reasoning with direct proofs over program heaps is explained in §3.

Definition 2 (Graph). A (partial) graph $G = (N, e)$ consists of a finite set of nodes $N \subseteq \mathfrak{N}$ and a mapping from pairs of nodes to edge functions $e: N \times \mathfrak{N} \rightarrow E$.

Flow Values and Flows. Flow values (taken from M ; the first element of a flow domain) are used to capture sufficient information to express desired non-local properties of a graph. In Example 1, flow values are non-negative integers; for the PIP (Example 2) we instead use *multisets* of integers, representing relevant *non-local* information: the priorities of nodes currently referencing a given node in the graph. Given such flow values, a node’s correct priority can be defined locally per node in the graph. This definition requires only the *maximum* value of these multisets, but as we will see shortly these multisets enable local *recomputation* of a correct priority when the graph is changed.

For a graph $G = (N, e)$ we express properties of G in terms of node-local conditions that may depend on the nodes’ *flow*. A flow is a function $fl: N \rightarrow M$ assigning every

node a flow value and must be some fixpoint of the following *flow equation*:

$$\forall n \in N. fl(n) = in(n) + \sum_{n' \in N} fl(n') \triangleright e(n', n) \quad (\text{FlowEqn})$$

Intuitively, one can think of the flow as being obtained by a fold computation over the graph:⁴ the *inflow* $in: N \rightarrow M$ defines an initial flow at each node. This initial flow is then updated recursively for each node n : the current flow value at its predecessor nodes n' is transferred to n via *edge functions* $e(n', n): M \rightarrow M$. These flow values are aggregated using the *summation operation* $+$ of the flow domain to obtain an updated flow of n ; a flow for the graph is some fixpoint satisfying this equation at all nodes.⁵

Definition 3 (Flow Graph). A flow graph $H = (N, e, fl)$ is a graph (N, e) and function $fl: N \rightarrow M$ such that there exists an inflow $in: N \rightarrow M$ satisfying FlowEqn(in, e, fl).

We let $\text{dom}(H) = N$, and sometimes identify H and $\text{dom}(H)$ to ease notational burden. For $n \in H$ we write H_n for the singleton flow subgraph of H induced by n .

Edge Functions. In any flow graph, the flow value assigned to a node n by a flow is propagated to its neighbours n' (and transitively) according to the edge function $e(n, n')$ labelling the edge (n, n') . The edge function maps the flow value at the *source node* n to one propagated on *this edge* to the *target node* n' . Note that we require such a labelling for *all* pairs consisting of a source node n inside the graph and a target node $n' \in \mathfrak{N}$ (i.e., possibly outside the graph). The 0 flow value (the third element of our flow domains) is used to represent no flow; the corresponding (constant) zero function $\lambda_0 = (\lambda m. 0)$ is used as edge function to model the *absence* of an edge in the graph. A set of edge functions E from which this labelling is chosen can, other than the requirement $\lambda_0 \in E$, be chosen as desired. As we will see in §4.4, restrictions to particular sets of edge functions E can be exploited to further strengthen our overall technique. Edge functions can depend on the local state of the source node (as in the following example); dependencies from elsewhere in the graph must be represented by the node's flow.

Example 3. Consider the graph in Figure 1 and the flow domain as in Example 2. We choose the edge functions to be λ_0 where no edge exists in the PIP structure, and otherwise $(\lambda m. \{\max(m \cup \{d\})\})$ where d is the default priority of the source of the edge. For example, in Figure 1, $e(r_3, p_2) = \lambda_0$ and $e(r_3, p_1) = (\lambda m. \{\max(m \cup \{0\})\})$. Since the flow value at r_3 is $\{1, 2, 2\}$, the edge (r_3, p_1) propagates the value $\{2\}$ to p_1 , correctly representing the current priority of r_3 .

Flow Aggregation and Inflows. The flow value at a node is defined by those propagated to it from each node in a graph via edge functions, along with an additional *inflow* value explained here. Since multiple non-zero flow values can be propagated to a node, we require an aggregation of these values via a binary $+$ operator on flow values: the second

⁴ We note that flows are not generally defined in this manner as we consider any fixpoint of the flow equation to be a flow. Nonetheless, the analogy helps to build an initial intuition.

⁵ We discuss questions regarding the existence and uniqueness of such fixpoints in §4.

element of our flow domains. The edges from which the aggregated values originate are unordered. Thus, we require $+$ to be commutative and associative, making this aggregation order-independent. The 0 flow value (**representing no flow**) must act as a unit for $+$. For example, in the path-counting flow domain $+$ means addition on natural numbers, while for the multisets employed for the PIP it means multiset union.

Each node in a flow graph has an *inflow*, modelling contributions to its flow value which do *not* come from inside the graph. Inflows play two important roles: first, since our graphs are partial, they model contributions from nodes *outside of the graph*. Second, inflow can be artificially added as a means of specialising the computation of flow values to characterise specific graph properties. For example, in the path-counting domain, we give an inflow of 1 to the node from which we are counting paths, and 0 to all others.

Example 4. Let the edges in the graph in Figure 1 be labelled as described in Example 3. If the inflow function in assigns the empty multiset to every node n and we let $fl(n)$ be the multiset labelling every node in the figure, then $\text{FlowEqn}(in, e, fl)$ holds.

The flow equation (FlowEqn) defines the flow of a node n to be the aggregation of flow values coming from other nodes n' inside the graph (as given by the respective edge function $e(n', n)$) as well as the inflow $in(n)$. Preserving solutions to this equation across updates to the graph structure is a fundamental goal of our technique. The following lemma (which relies on the fact that $+$ is required to be cancellative) states that any correct flow values uniquely determine appropriate inflow values:⁶

Lemma 1. *Given a flow graph (N, e, fl) , there exists a unique inflow in such that $\text{FlowEqn}(in, e, fl)$.*

We now turn to how solutions of the flow equation can be preserved or appropriately updated under *changes* to the underlying graph.

Graph Updates and Cancellativity. Given a flow graph with known flow and inflow values, suppose we *remove* an edge from n_1 to n_2 (replacing the edge function with λ_0). For the same inflow, such an update will potentially affect the flow at n_2 and nodes to which n_2 (transitively) propagates flow. Starting from the simple case that n_2 has no outgoing edges, we need to recompute a suitable flow at n_2 . Knowing the old flow value (say, m) and the contribution $m' = fl(n_1) \triangleright e(n_1, n_2)$ *previously* provided along the removed edge, we know that the correct new flow value is some m'' such that $m' + m'' = m$. This constraint has a unique solution (and thus, we can unambiguously recompute a new flow value) exactly when the aggregation $+$ is *cancellative*; we therefore make cancellativity a *requirement* on the $+$ of any flow domain.

Cancellativity intuitively enforces that the flow domain carries enough information to enable adaptation to local updates (in particular, removal of edges⁷). Returning to the PIP example, cancellativity requires us to carry multisets as flow values rather than only the maximum priority value: $+$ cannot be the maximum operation, as this would not be

⁶ Proofs for all lemmas and theorems can be found in Appendix E.

⁷ As we will show in §2.3, an analogous problem for composition of flow graphs is also directly solved by this choice to force aggregation to be cancellative.

cancellative. The resulting multisets (like the `prio` fields in the actual code) provide the information necessary to recompute corrected priority values locally.

For example, in the PIP graph shown in Figure 1, removing the edge from p_6 to r_4 would not affect the current priority of r_4 whereas if p_7 had current priority 1 instead of 2, then the current priority of r_4 would have to decrease. In either case, recomputing the flow value for r_4 is simply a matter of subtraction (removing $\{2\}$ from the multiset at r_4); cancellativity guarantees that our flow domains will always provide the information needed for this recomputation. Without this property, the recomputation of a flow value for the target node n_2 would, in general, entail recomputing the incoming flow values from all remaining edges from scratch. Cancellativity is also crucial for Lemma 1 above, forcing uniqueness of inflows, given known flow values in a flow graph. This allows us to define natural but powerful notions of flow graph decomposition and recomposition.

2.3 Flow Graph Composition and Abstraction

Building towards the core of our reasoning technique, we now turn to the question of decomposition and recomposition of flow graphs. Two flow graphs with disjoint domains always compose to a graph, but this will be a *flow graph* only if their flows are chosen consistently to admit a solution to the resulting flow equation (i.e. the flow graph composition operator \odot defined below is *partial*).

Definition 4 (Flow Graph Algebra). *The flow graph algebra (FG, \odot, H_\emptyset) for the flow domain $(M, +, 0, E)$ is defined by*

$$\begin{aligned} FG &:= \{(N, e, fl) \mid (N, e, fl) \text{ is a flow graph}\}, & H_\emptyset &:= (\emptyset, e_\emptyset, fl_\emptyset), \\ (N_1, e_1, fl_1) \odot (N_2, e_2, fl_2) &:= \begin{cases} (N_1 \uplus N_2, e_1 \uplus e_2, fl_1 \uplus fl_2) & \text{if } inFG \\ \perp & \text{otherwise,} \end{cases} \end{aligned}$$

where e_\emptyset and fl_\emptyset are the edge functions and flow on the empty set of nodes $N = \emptyset$.

Intuitively, two flow graphs compose to a flow graph if their contributions to each others' flow (along edges from one to the other) are reflected in the corresponding inflow of the other graph. For example, consider the subgraph from Figure 1 consisting of the single node p_7 (with 0 inflow). This will compose with the remainder of the graph depicted only if this remainder subgraph has an inflow which, at node r_4 , includes at least the multiset $\{2\}$, reflecting the propagated value from p_7 .

We use this intuition to extract an *abstraction* of flow graphs which we call *flow interfaces*. Given a flow (sub)graph, its flow interface consists of the node-wise inflow and *outflow* (the flow contributions its nodes make to all nodes outside of the graph, defined below). It is thus an abstraction that hides the flow values and edges that are wholly *inside* the flow graph. Flow graphs that have the same flow interface “look the same” to the external graph, as the same values are propagated inwards and outwards.

Definition 5 (Flow Interface). *For a given flow domain M , a flow interface is a pair $I = (in, out)$ where $in: N \rightarrow M$ and $out: \mathfrak{N} \setminus N \rightarrow M$ for some $N \subseteq \mathfrak{N}$.*

We write $I.in, I.out$ for the two components of the interface $I = (in, out)$. We will again sometimes identify I and $\text{dom}(I.in)$ to ease notational burden.

Given a flow graph $H \in \text{FG}$, we can compute its interface as follows. Recall that Lemma 1 implies that any flow graph has a unique inflow. Thus, we can define an inflow function that maps each flow graph $H = (N, e, fl)$ to the unique inflow $\text{inf}(H): H \rightarrow M$ such that $\text{FlowEqn}(\text{inf}(H), e, fl)$. Dually, we define the *outflow* of H as the function $\text{outf}(H): \mathfrak{N} \setminus N \rightarrow M$ defined by $\text{outf}(H)(n) := \sum_{n' \in N} fl(n') \triangleright e(n', n)$. The *flow interface of H* , written $\text{int}(H)$, is the pair $(\text{inf}(H), \text{outf}(H))$ consisting of its inflow and its outflow. Returning to the previous example, if H is the singleton subgraph consisting of node p_7 from Figure 1 with flow and edges as depicted, then $\text{int}(H) = (\lambda n. \emptyset, \lambda n. (n=r_4 ? \{2\} : \emptyset))$.

This abstraction, while simple, turns out to be powerful enough to build a separation algebra over our flow graphs, allowing them to be decomposed, locally modified and recomposed in ways yielding all the local reasoning benefits of separation logics. In particular, for graph operations within a subgraph with a certain interface, we need to prove: (a) that the modified subgraph is still a flow graph (by checking that the flow equation still has a solution locally in the subgraph) and (b) that it satisfies the same interface (in other words, the effect of the modification on the flow is contained within the subgraph); the meta-level results for our technique then justify that we can recombine the modified subgraph with any graph that the original could be composed with.

We define the corresponding *flow interface algebra* as follows:

Definition 6 (Flow Interface Algebra). For a given flow domain M , the flow interface algebra over M is defined to be $(\text{FI}, \oplus, I_\emptyset)$, where:

$$\begin{aligned} \text{FI} &:= \{I \mid I \text{ is a flow interface}\}, & I_\emptyset &:= \text{int}(H_\emptyset), \\ I_1 \oplus I_2 &:= \begin{cases} I & I_1 \cap I_2 = \emptyset \\ & \wedge \forall i \neq j \in \{1, 2\}, n \in I_i. I_i.in(n) = I.in(n) + I_j.out(n) \\ & \wedge \forall n \notin I. I.out(n) = I_1.out(n) + I_2.out(n) \\ \perp & \text{otherwise.} \end{cases} \end{aligned}$$

Flow interface composition is well-defined because of cancellativity of the underlying flow domain (it is also, exactly as flow graph composition, partial).

We next show the key result for this abstraction: the ability for two flow graphs to compose depends only on their interfaces; flow interfaces implicitly define a congruence relation on flow graphs.

Lemma 2. $\text{int}(H_1) = I_1 \wedge \text{int}(H_2) = I_2 \Rightarrow \text{int}(H_1 \odot H_2) = I_1 \oplus I_2$.

Crucially, the following result shows that we can use our flow interfaces as an abstraction directly compatible with existing separation logics.

Theorem 1. *The flow interface algebra $(\text{FI}, \oplus, I_\emptyset)$ is a separation algebra.*

This result forms the core of our reasoning technique; it enables us to make modifications within a chosen subgraph and, by proving preservation of its interface, know that the result composes with any context exactly as the original did. Flow interfaces capture

precisely the information relevant about a flow graph, with respect to composition with other flow graphs. In [Appendix B](#) we provide additional examples of flow domains that demonstrate the range of data structures and graph properties that can be expressed using flows, including a notion of *universal flow* that in a sense provides a completeness result for the expressivity of the framework. We now turn to constructing proofs atop these new reasoning principles.

3 Proof Technique

This section shows how to integrate flow reasoning into a standard separation logic, using the priority inheritance protocol (PIP) algorithm to illustrate our proof techniques.

Since flow graphs and flow interfaces form separation algebras, it is possible in principle to define a separation logic (SL) using these notions as a custom *semantic model* (indeed, this is the proof approach taken in [25]). By contrast, we integrate flow interfaces with a *standard* separation logic without modifying its semantics. This has the important technical advantage that our proof technique can be naturally integrated with existing separation logics and verification tools supporting SL-style reasoning. We consider a standard *sequential* SL in this section, but our technique can also be directly integrated with a concurrent SL such as RGSep (as we show in §4.5) or frameworks such as Iris [21] supporting (ghost) resources ranging over user-defined separation algebras.

3.1 Encoding Flow-based Proofs in SL

Proofs using our flow framework can employ a combination of specifications enforced at the node level and in terms of the flow graphs and interfaces corresponding to larger heap regions such as entire data structures (henceforth, *composite graphs* and *composite interfaces*). At the node level, we write invariants that every node is intended to satisfy, typically relating the node’s flow value to its local state (fields). For example, in the PIP, we use node-local invariants to express that a node’s current priority is the maximum of the node’s default priority and those in its current flow value. We typically express such specifications in terms of *singleton (flow) graphs*, and their *singleton interfaces*.

Specification in terms of *composite* interfaces has several important purposes. One is to define custom inflows: e.g. in the path-counting flow domain, specifying that the inflow of a composite interface is 1 at some designated node r and 0 elsewhere enforces in any underlying flow graph that each node n ’s flow value will be the number of paths from r to n .⁸ Composite interfaces can also be used to express that, in two states of execution, a portion of the heap “looks the same” with respect to composition (it has the same interface, and so can be composed with the same flow graphs), or to capture by *how much* there is an observable difference in inflow or outflow; we employ this idea in the PIP proof below.

We now define an assertion syntax convenient for capturing both node-level and composite-level constraints, defined within an SL-style proof system. We assume an *intu-*

⁸ Note that the analogous property cannot be captured at the node level; when considering singleton interfaces per node in a tree rooted at r , *every* singleton interface has an inflow of 1.

tionistic, garbage-collected SL [7] with standard syntax and semantics:⁹ see [Appendix A](#) for more details.

Node Predicates. The basic building block of our flow-based specifications is a node predicate $N(x, H)$, representing ownership of the fields of a single node x , as well as capturing its corresponding singleton flow graph H :

$$N(x, H) := \exists fs, fl. x \mapsto fs * H = (\{x\}, (\lambda y. \text{edge}(x, fs, y)), fl) * \gamma(x, fs, fl(x))$$

N is implicitly parameterised by fs , edge and γ ; these are explained next and are typically fixed across any given flow-based proof. The N predicate expresses that we have a heap cell at location x containing fields fs (a list of field-name/value mappings).¹⁰ It also says that H is a singleton flow graph with domain $\{x\}$ with some flow fl , whose edge functions are defined by a user-defined abstraction function $\text{edge}(x, fs, y)$; this function allows us to define edges in terms of x 's field values. Finally, the node, its fields, and its flow in this flow graph satisfy the custom predicate γ , used to encode node-local properties such as constraints in terms of the flow values of nodes.

Graph Predicates. The analogous predicate for composite graphs is Gr . It carries ownership to the nodes making up a potentially unbounded graph, using iterated separating conjunction over a *set* of nodes X as mentioned in §1:

$$Gr(X, H) := \exists \mathcal{H}. \bigstar_{x \in X} N(x, \mathcal{H}(x)) * H = \bigodot_{x \in X} \mathcal{H}(x)$$

Gr is also implicitly parameterised by fs , edge and γ . The existentially-quantified \mathcal{H} is a logical variable representing a *function* from nodes in X to corresponding singleton flow graphs. $Gr(X, H)$ describes a set of nodes X , such that each $x \in X$ is an N (in particular, it satisfies γ), whose singleton flow graphs compose back to H . As well as carrying ownership of the underlying heap locations, Gr 's definition allows us to connect a node-level view of the region X (each $\mathcal{H}(x)$) with a composite-level view defined by H , on which we can impose appropriate graph-level properties such as constraints on the region's inflow.

Lifting to Interfaces. Flow based proofs can often be expressed more elegantly and abstractly using predicates in terms of node and composite-level interfaces rather than flow graphs. To this end, we overload both our node and graph predicates with analogues whose second parameter is a flow interface, defined as follows:

$$\begin{aligned} N(x, I) &:= \exists H. N(x, H) * I = \text{int}(H) \\ Gr(X, I) &:= \exists H. Gr(x, H) * I = \text{int}(H) \end{aligned}$$

We will use these versions in the PIP proof below; interfaces capture all relevant properties for decomposition and composition of these flow graphs.

$$\begin{array}{lcl}
\text{Gr}(X_1 \uplus X_2, H) & \models & \exists H_1, H_2. \text{Gr}(X_1, H_1) * \text{Gr}(X_2, H_2) \\
& & * H_1 \odot H_2 = H \quad (\text{DECOMP}) \\
\text{Gr}(X_1, H_1) * \text{Gr}(X_2, H_2) * H_1 \odot H_2 \neq \perp & \models & \text{Gr}(X_1 \uplus X_2, H_1 \odot H_2) \quad (\text{COMP}) \\
\text{N}(x, H) & \equiv & \text{Gr}(\{x\}, H) \quad (\text{SING}) \\
\text{emp} & \models & \text{Gr}(\emptyset, H_\emptyset) \quad (\text{GEMP}) \\
\text{Gr}(X_1, H'_1) * \text{Gr}(X_2, H_2) * H = H_1 \odot H_2 & \models & \text{Gr}(X_1 \uplus X_2, H'_1 \odot H_2) \quad (\text{REPL}) \\
* \text{int}(H_1) = \text{int}(H'_1) & & * \text{int}(H) = \text{int}(H'_1 \odot H_2)
\end{array}$$

Fig. 2: Some useful lemmas for proving entailments between flow-based specifications.

Flow Lemmas. We first illustrate our N and Gr predicates (which capture SL ownership of heap regions and abstract these with flow interfaces) by identifying a number of lemmas which are generically useful in flow-based proofs. Reasoning at the level of flow interfaces is entirely in the *pure* world (mathematics independent of heap-ownership and resources) with respect to the underlying SL reasoning; these lemmas are consequences of our predicate definitions and the foundational flow framework definitions themselves.

Examples of these lemmas are shown in Figure 2. (DECOMP) shows that we can always decompose a valid flow graph into subgraphs which are themselves flow graphs. Recomposition (COMP) is possible only if the subgraphs compose. These rules, as well as (SING), and (GEMP) follow directly from the definition of Gr and standard SL properties of iterated separating conjunction. The final rule (REPL) is a direct consequence of rules (COMP), (DECOMP) and the congruence relation on flow graphs induced by their interfaces (cf. Lemma 2). Conceptually, it expresses that after decomposing any flow graph into two parts H_1 and H_2 , we can *replace* H_1 with a new flow graph H'_1 with the same interface; when recomposing, the overall graph will be a flow graph with the same overall interface.

Note the connection between rules (COMP)/(DECOMP) and the algebraic laws of standard inductive predicates such as ls describing a segment of a linked list [2]. For instance by combining the definition of Gr with these rules and (SING) we can prove the following graph analogue of the rule to separate a list into the head node and the tail:

$$\text{Gr}(X \uplus \{y\}, H) \equiv \exists H_y, H'. \text{N}(y, H_y) * \text{Gr}(X, H') * H = H_y \odot H' \quad ((\text{UN})\text{FOLD})$$

However, crucially (and unlike when using general inductive predicates [35]), this rule is symmetrical for any node x in X ; it works analogously for any desired order of decomposition of the graph, and for any data structure specified using flows.

When working with our overloaded N and Gr predicates, similar steps to those described by the above lemmas are useful. Given these overloaded predicates, we simply

⁹ As $P * \phi \equiv P \wedge \phi$ for pure formulas P in garbage-collected SLs, we use $*$ instead of \wedge throughout this paper.

¹⁰ For simplicity, we assume that all fields of a flow graph node are to be handled by our flow-based technique, and that their ownership (via \mapsto points-to predicates) is always carried around together; lifting these restrictions would be straightforward.

apply the lemmas above to the *existentially quantified* flow-graphs in their definitions and then lift the consequence of the lemma back to the interface level using the congruence between our flow graph and interface composition notions (Lemma 2).

3.2 Proof of the PIP

We now have all the tools necessary to verify the priority inheritance protocol (PIP). Figure 3 gives the full algorithm with flow-based specifications; we also include some intermediate assertions to illustrate the reasoning steps for the `acquire` method, which we explain in more detail below.¹¹ We instantiate our framework in order to capture the PIP invariants as follows:

$$\begin{aligned}
 fs &:= \{\text{next}: y, \text{curr_prio}: q, \text{def_prio}: q^0, \text{prios}: Q\} \\
 \text{edge}(x, fs, z) &:= \begin{cases} (\lambda m. \max(m \cup \{q^0\})) & \text{if } z = y \neq \text{null} \\ \lambda_0 & \text{otherwise} \end{cases} \\
 \gamma(x, fs, m) &:= q^0 \geq 0 * (\forall q' \in Q. q' \geq 0) * m = Q * q = \{\max(Q \cup \{q^0\})\} \\
 \varphi(I) &:= I = (\lambda_0, \lambda_0)
 \end{aligned}$$

Each node has the four fields listed in *fs*. *fs* also defines variables such as *y* to denote field values that are used in the definitions of *edge* and γ ; these variables are bound to the heap by `N`. *edge* abstracts the heap into a flow graph by letting each node have an edge to its `next` successor labelled by a function that passes to it the maximum incoming priority or the node’s default priority: whichever is larger. With this definition, one can see that the flow of every node will be the multiset containing exactly the priorities of its predecessors. The node-local invariant γ says that all priorities are non-negative, the flow *m* of each node is stored in the `prios` field, and its current priority is the maximum of its default and incoming priorities. Finally, the constraint φ on the global interface expresses that the graph is closed – it has no inflow or outflow.

Flows Specifications for the PIP. Our specifications of `acquire` and `release` guarantee that if we start with a valid flow graph (closed, according to φ), we are guaranteed to return a valid flow graph with the same interface (i.e. the graph remains closed). For clarity of the exposition, we focus here on how we prove that being a flow graph that satisfies the PIP invariant is preserved (as is the composite flow graph’s interface). Extending this specification to one which proves, e.g., that `acquire` adds the expected edge is straightforward (see [Appendix C](#)).¹²

The specification for `update` is somewhat subtle, and exploits the full flexibility of flow interfaces as a specification medium. The preconditions of `update` describe an update to the graph which is not yet completed. There are three complementary aspects to this specification. Firstly, (as for `acquire` and `release`), node-local invariants (γ) hold for all nodes in the graph (enforced via `N` and `Gr` predicates). Secondly, we employ

¹¹ In specifications, we implicitly quantify at the top level over free variables such as *I*. λ_0 denotes an identically zero function on an unconstrained domain.

¹² We also omit `acquire`’s precondition that `p.next == null` for brevity.

```

1 // Let  $\delta(m, q_1, q_2) := m \setminus (q_1 \geq 0 ? \{q_1\} : \emptyset) \cup (q_2 \geq 0 ? \{q_2\} : \emptyset)$ 
2
3 method update(n: Ref, from: Int, to: Int)
4   requires  $\mathbf{N}(n, I_n) * \mathbf{Gr}(X \setminus \{n\}, I') * I = I'_n \oplus I' * \varphi(I) * n \in X$ 
5   requires  $I'_n = (\{n \mapsto \delta(I_n.in(n), from, to)\}, I_n.out) * from \neq to$ 
6   ensures  $\mathbf{Gr}(X, I)$ 
7 {
8   n.prios := n.prios \ {from}
9   if (to >= 0) {
10    n.prios := n.prios  $\cup$  {to}
11  }
12  from := n.curr_prio
13  n.curr_prio := max(n.prios  $\cup$  {n.def_prio})
14  to := n.curr_prio
15
16  if (from != to && n.next != null) {
17    update(n.next, from, to)
18  }
19 }
20
21 method acquire(p: Ref, r: Ref)
22   requires  $\mathbf{Gr}(X, I) * \varphi(I) * p \in X * r \in X * p \neq r$ 
23   ensures  $\mathbf{Gr}(X, I)$ 
24 {
25    $\{\exists I_r, I_p, I_1. \mathbf{N}(r, I_r) * \mathbf{N}(p, I_p) * \mathbf{Gr}(X \setminus \{r, p\}, I_1) * I = I_r \oplus I_p \oplus I_1 * \varphi(I)\}$ 
26   if (r.next == null) {
27     r.next := p;
28     // Let  $q_r = r.curr\_prio$ 
29      $\left\{ \begin{array}{l} \exists I_r, I'_r, I_p, I_1. \mathbf{N}(r, I'_r) * \mathbf{N}(p, I_p) * \mathbf{Gr}(X \setminus \{r, p\}, I_1) * I = I_r \oplus I_p \oplus I_1 \\ * I'_r = (I_r.in, \{p \mapsto \{q_r\}\}) * I_r.out = \lambda_0 * \dots \end{array} \right\}$ 
30      $\equiv \left\{ \begin{array}{l} \exists I_p, I'_p, I_2. \mathbf{N}(p, I_p) * \mathbf{Gr}(X \setminus \{p\}, I_2) * I = I'_p \oplus I_2 \\ * I'_p = (\{p \mapsto \delta(I_p.in(p), -1, q_r)\}, I_p.out) * \dots \end{array} \right\}$ 
31     update(p, -1, r.curr_prio)
32      $\{\mathbf{Gr}(X, I)\}$ 
33   } else {
34     p.next := r; update(r, -1, p.curr_prio)
35   }
36 }
37
38 method release(p: Ref, r: Ref)
39   requires  $\mathbf{Gr}(X, I) * \varphi(I) * p \in X * r \in X * p \neq r$ 
40   ensures  $\mathbf{Gr}(X, I)$ 
41 { r.next := null; update(p, r.curr_prio, -1) }

```

Fig. 3: Full PIP code and specifications, with proof sketch for `acquire`. The comments and coloured annotations (lines 29 to 32) are used to highlight steps in the proof, and are explained in detail in the text.

flow interfaces to express a decomposition of the original top-level interface I into compatible (primed) sub-interfaces. The key to understanding this specification is that I'_n is in some sense a *fake* interface; it does not abstract the current state of the heap node n . Instead, I'_n expresses the way in which the node n 's current inflow *hasn't yet* been accounted for in the heap: that *if* n could adjust its inflow according to the propagated priority change *without* changing its outflow, then it would compose back with the rest of the graph, and restore the graph's overall interface. The shorthand δ defines the required change to n 's inflow.

In general (except when n 's `next` field is null, or n 's flow value is unchanged), it is not even possible for n 's fields to be updated to satisfy I'_n ; by updating n 's inflow, we will necessarily update its outflow. However, we can then construct a corresponding “fake” interface for the next node in the graph, reflecting the update yet to be accounted for, and establishing the precondition for the recursive call to `update`.

The third specification aspect is the *connection* between heap-level nodes and interfaces. The $N(n, I_n)$ predicate connects n with a *different* interface; I_n is the actual current abstraction of n 's state. Conceptually, the key property which is broken at this point is this connection between the interface-level specification and the heap at node n , reflected by the decomposition in the specification between $X \setminus \{n\}$ and $\{n\}$.

We note that the same specification ideas and proof style can be easily adapted to other data structure implementations with an update-notify style, including well-known designs such as Subject-Observer patterns, or the Composite pattern [29].

Proof Outline. To illustrate the application of flows reasoning to our PIP specification ideas more clearly, we examine in detail the first `if`-branch in the proof of `acquire`. Our intermediate proof steps are shown as purple annotations surrounded by braces. The first step, as shown in the first line inside the method body, is to apply ((UN)FOLD) twice (on the flow graphs represented by these predicates) and peel off N predicates for each of r and p . The update to r 's `next` field (line 27) causes the correct singleton interface of r to change to I'_r : its outflow (previously none, since the `next` field was null) now propagates flow to p . We summarise this state in the assertion on line 29 (we omit e.g. repetition of properties from the function's precondition, focusing on the flow-related steps of the argument). We now rewrite this state; using the definition of interface composition (Definition 6) we deduce that although I'_r and I_p do not compose (since the former has outflow that the latter does not account for as inflow), the alternative “fake” interface I'_p for p (which artificially accounts for the missing inflow) *would* do so (cf. line 30). Essentially, we show $I_r \oplus I_p = I'_r \oplus I'_p$, that the interface of $\{r, p\}$ would be unchanged if p could somehow have interface I'_p . Now by setting $I_2 = I'_r \oplus I_1$ and using algebraic properties of interfaces, we assemble the precondition expected by `update`. After the call, `update`'s postcondition gives us the desired postcondition.

We focused here on the details of `acquire`'s proof, but very similar manipulations are required for reasoning about the recursive call in `update`'s implementation.¹³ The main difference there is that if the `if`-condition wrapping the recursive call is false then either the last-modified node has no successor (and so there is no outstanding inflow

¹³ We provide further proof outlines in [Appendix C](#).




Fig. 4: A potential state of the Harris list with explicit memory management. f_{next} pointers are shown with dashed edges, marked nodes are shaded gray, and null pointers are omitted for clarity.

change needed), or we have $f_{from} = t_o$ which implies that the “fake” interface is actually the same as the currently correct one.

Despite the property proved for the PIP example being a rather delicate recursive invariant over the (potentially cyclic) graph, the power of our framework enables extremely succinct specifications for the example, and proofs which require the application of relatively few generic lemmas. The integration with standard separation logic reasoning, and the complementary separation algebras provided by flow interfaces allow decomposition and recomposition to be simple proof steps. For this proof, we integrated with standard sequential separation logic, but in the next section we will show that compatibility with concurrent SL techniques is similarly straightforward.

4 Advanced Flow Reasoning and the Harris List

This section introduces some advanced foundational flow framework theory and demonstrates its use in the proof of the Harris list. We note that [25] presented a proof of this data structure in the original flow framework. The proof given here shows that the new framework eliminates the need for the customized concurrent separation logic defined in [25]. We start with a recap of Harris’ algorithm adapted from [25].

4.1 The Harris List Algorithm

The power of flow-based reasoning is exhibited in the proof of overlaid data structures such as the Harris list, a concurrent non-blocking linked list algorithm [13]. This algorithm implements a set data structure as a sorted list, and uses atomic compare-and-swap (CAS) operations to allow a high degree of parallelism. As with the sequential linked list, Harris’ algorithm inserts a new key k into the list by finding nodes k_1, k_2 such that $k_1 < k < k_2$, setting k to point to k_2 , and using a CAS to change k_1 to point to k only if it was still pointing to k_2 . However, a similar approach fails for the delete operation. If we had consecutive nodes k_1, k_2, k_3 and we wanted to delete k_2 from the list (say by setting k_1 to point to k_3), there is no way to ensure with one CAS that k_2 and k_3 are also still adjacent (another thread could have inserted/deleted in between them).

Harris’ solution is a two step deletion: first atomically mark k_2 as deleted (by setting a mark bit on its successor field) and then later remove it from the list using a single

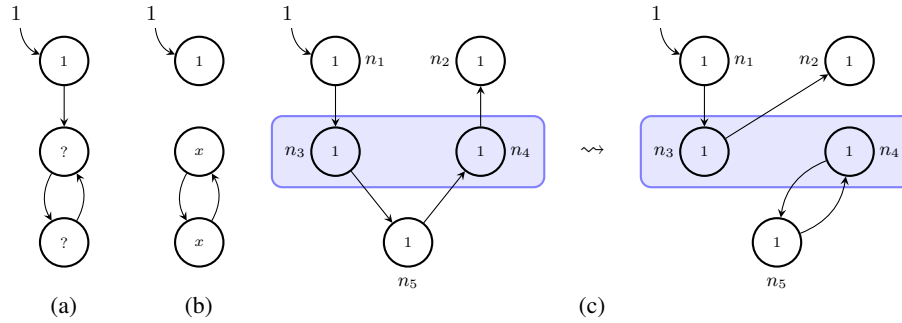


Fig. 5: Examples of graphs that motivate effective acyclicity. All graphs use the path-counting flow domain, the flow is displayed inside each node, and the inflow is displayed as curved arrows to the top-left of nodes. (a) shows a graph and inflow that has no solution to (FlowEqn); (b) has many solutions. (c) shows a modification that preserves the interface of the modified nodes, yet goes from a graph that has a unique flow to one that has many solutions to (FlowEqn).

CAS. After a node is marked, no thread can insert or delete to its right, hence a thread that wanted to insert k' to the right of k_2 would first remove k_2 from the list and then insert k' as the successor of k_1 .

In a non-garbage-collected environment, unlinked nodes cannot be immediately freed as suspended threads might continue to hold a reference to them. A common solution is to maintain a second “free list” to which marked nodes are added before they are unlinked from the main list (this is the so-called drain technique). These nodes are then labelled with a timestamp, which is used by a maintenance thread to free them when it is safe to do so. This leads to the kind of data structure shown in Figure 4, where each node has two pointer fields: a `next` field for the main list and an `fnext` field for the free list (the list from fh to ft via dashed edges). Threads that have been suspended while holding a reference to a node that was added to the free list can simply continue traversing the `next` pointers to find their way back to the unmarked nodes of the main list.

Even for seemingly simple properties such as that the Harris list is memory safe and not leaking memory, the proof will rely on the following non-trivial invariants:

- (a) The data structure consists of two (potentially overlapping) lists: a list on `next` edges beginning at mh and one on `fnext` edges beginning at fh .
- (b) The two lists are null terminated and `next` edges from nodes in the free list point to nodes in the free list or main list.
- (c) All nodes in the free list are marked.
- (d) ft is an element in the free list (due to concurrency, it’s not always the tail).

Challenges. To prove that Harris’ algorithm maintains the invariants listed above we must tackle a number of challenges. First, we must construct flow domains that allow us to describe overlaid data structures, such as the overlapping main and free lists (§4.2). Second, the flow-based proofs we have seen so far work by showing that the interface of

some modified region is unchanged. However, if we consider a program that allocates and inserts a new node into a data structure (like the insert method of Harris), then the interface cannot be the same since the domain has changed (it has increased by the newly allocated node). We must thus have a means to reason about preservation of flows by modifications that allocate new nodes (§4.3). The third issue is that in some flow domains, there exist graphs G and inflows in for which no solutions to the flow equation (FlowEqn) exist. For instance, consider the path-counting flow domain and the graph in Figure 5(a). Since we would need to use the path-counting flow in the proof of the Harris list to encode its structural invariants, this presents a challenge (§4.4).

We will next see how to overcome these three challenges in turn, and then apply those solution to the proof of the Harris list in §4.5.

4.2 Product Flows for Reasoning about Overlays

An important fact about flows is that any flow of a graph over a product of two flow domains is the product of the flows on each flow domain component.

Lemma 3. *Given two flow domains $(M_1, +_1, 0_1, E_1)$ and $(M_2, +_2, 0_2, E_2)$, the product domain $(M_1 \times M_2, +, (0_1, 0_2), E)$ is a flow domain, where $+$ and E are the pointwise liftings of $(+_1, +_2)$ and (E_1, E_2) , respectively, to the domain $M_1 \times M_2$.*

This lemma greatly simplifies reasoning about overlaid graph structures; we will use the product of two path-counting flows to describe a structure consisting of two overlaid lists that make up the Harris list.

4.3 Contextual Extensions and the Replacement Theorem

In general, when modifying a flow graph H to another flow graph H' , requiring that H' satisfies *precisely* the same interface $\text{int}(H)$ can be too strong a condition as it does not permit allocating new nodes. Instead, we want to allow $\text{int}(H')$ to differ from $\text{int}(H)$ in that the new interface could have a larger domain, as long as the edges from the new nodes do not change the outflow of the modified region.

Definition 7. *An interface $I = (in, out)$ is contextually extended by $I' = (in', out')$, written $I \lesssim I'$, if and only if the following conditions all hold:*

- (1) $\text{dom}(in) \subseteq \text{dom}(in')$,
- (2) $\forall n \in \text{dom}(in). in(n) = in'(n)$, and
- (3) $\forall n' \notin \text{dom}(in'). out(n') = out'(n')$.

The following theorem states that contextual extension preserves composability and is itself preserved under interface composition.

Theorem 2 (Replacement Theorem). *If $I = I_1 \oplus I_2$, and $I_1 \lesssim I'_1$ are all valid interfaces such that $I'_1 \cap I_2 = \emptyset$ and $\forall n \in I'_1 \setminus I_1. I_2.out(n) = 0$, then there exists a valid $I' = I'_1 \oplus I_2$ such that $I \lesssim I'$.*

In terms of our flow predicates, this theorem gives rise to the following adaptation of the (REPL) rule:

$$\begin{aligned} & \text{Gr}(X'_1, H'_1) * \text{Gr}(X_2, H_2) * H = H_1 \odot H_2 * \text{int}(H_1) \lesssim \text{int}(H'_1) \\ \models & \exists H'. \text{Gr}(X'_1 \uplus X_2, H') * H' = H'_1 \odot H_2 * \text{int}(H) \lesssim \text{int}(H') \quad (\text{REPL+}) \end{aligned}$$

The rule (REPL+) is derived from the Replacement Theorem by instantiating with $I = \text{int}(H)$, $I_1 = \text{int}(H_1)$, $I_2 = \text{int}(H_2)$ and $I'_1 = \text{int}(H'_1)$. We know $I_1 \lesssim I'_1$; $H = H_1 \odot H_2$ tells us (by Lemma 2) that $I = I_1 \oplus I_2$, and $\text{Gr}(X'_1, H'_1) * \text{Gr}(X_2, H_2)$ gives us $I'_1 \cap I_2 = \emptyset$. The final condition of the Replacement Theorem is to prove that there is no outflow from X_2 to any newly allocated node in X'_1 . While we can use additional ghost state to prove such constraints in our proofs, if we assume that the memory allocator only allocates fresh addresses and restrict the abstraction function edge to only propagate flow along an edge (n, n') if n has a (non-ghost) field with a reference to n' then this condition is always true. For simplicity, and to keep the focus of this paper on the flow reasoning, we make this assumption in the Harris list proof.

4.4 Existence and Uniqueness of Flows

We typically express global properties of a graph $G = (N, e)$ by fixing a global inflow $in: N \rightarrow M$ and then constraining the flow of each node in N using node-local conditions. However, as we discussed at the beginning of this section, there is no general guarantee that a flow exists or is unique for a given in and G . The remainder of this section presents two complementary conditions under which we can prove that our flow fixpoint equation always has a unique solution. To this end, we say that a flow domain $(M, +, 0, E)$ has *unique flows* if for every graph (N, e) over this flow domain and inflow $in: N \rightarrow M$, there exists a unique fl that satisfies the flow equation $\text{FlowEqn}(in, e, fl)$. But first, we briefly recall some more monoid theory.

We say M is *positive* if $m_1 + m_2 = 0$ implies that $m_1 = m_2 = 0$. For a positive monoid M , we can define a partial order \leq on its elements as $m_1 \leq m_2$ if and only if $\exists m_3. m_1 + m_3 = m_2$. This definition implies that every $m \in M$ satisfies $0 \leq m$.

For $e, e': M \rightarrow M$, we write $e + e'$ for the function that maps $m \in M$ to $e(m) + e'(m)$. We lift this construction to a set of functions E and write it as $\sum_{e \in E} e$.

Definition 8. A function $e: M \rightarrow M$ is called an *endomorphism on M* if for every $m_1, m_2 \in M$, $e(m_1 + m_2) = e(m_1) + e(m_2)$. We denote the set of all endomorphisms on M by $\text{End}(M)$.

Note that for cancellative M , $e(0) = 0$ for every endomorphism $e \in \text{End}(M)$. Note further that $e + e' \in \text{End}(M)$ for any $e, e' \in \text{End}(M)$. Similarly, for finite sets $E \subseteq \text{End}(M)$, $\sum_{e \in E} e \in \text{End}(M)$. We say that a set of endomorphisms $E \subseteq \text{End}(M)$ is *closed* if for every $e, e' \in E$, $e \circ e' \in E$ and $e + e' \in E$.

Nilpotent Cycles. Let $(M, +, 0, E)$ be a flow domain where every edge function $e \in E$ is an endomorphism on M . In this case, we can show that the flow of a node n is the sum of the flow as computed along *each path* in the graph that ends at n . Suppose we additionally know that the edge functions are defined such that their composition along any *cycle* in the graph eventually becomes the identically zero function. We then need only consider finitely many paths to compute the flow of a node, which means the flow equation has a unique solution.

Formally, such edge functions are called *nilpotent endomorphisms*:

Definition 9. A closed set of endomorphisms $E \subseteq \text{End}(M)$ is called *nilpotent* if there exists $p > 1$ such that $e^p \equiv 0$ for every $e \in E$.

Example 5. The flow domain $(\mathbb{N}^2, +, (0, 0), \{(\lambda(x, y), (0, c \cdot x)) \mid c \in \mathbb{N}\})$ contains nilpotent edge functions that shift the first component of the flow to the second (with a scaling factor). This domain can be used to express the property that every node in a graph is reachable from the root via a single edge (by requiring the flow of every node to be $(0, 1)$ under the inflow $(\lambda n. (n = r ? (1, 0) : (0, 0)))$).

Before we prove that nilpotent endomorphisms lead to unique flows, we present a useful notion when dealing with endomorphic flow domains.

Definition 10. The capacity of a flow graph $G = (N, e)$ is $\text{cap}(G): N \times \mathfrak{N} \rightarrow (M \rightarrow M)$, defined inductively as $\text{cap}(G) := \text{cap}^{|G|}(G)$, where $\text{cap}^0(G)(n, n') := \delta_{n=n'}$ and

$$\text{cap}^{i+1}(G)(n, n') := \delta_{n=n'} + \sum_{n'' \in G} \text{cap}^i(G)(n, n'') \circ e(n'', n').$$

For a flow graph $H = (N, e, fl)$, we write $\text{cap}(H)(n, n') = \text{cap}((N, e))(n, n')$ for the capacity of the underlying graph. Intuitively, $\text{cap}(G)(n, n')$ is the function that summarizes how flow is routed from any source node n in G to any other node n' , including those outside of G .

We can now show that if all edges of a flow graph are labelled with edges from a nilpotent set of endomorphisms, then the flow equation has a unique solution:

Lemma 4. If $(M, +, 0, E)$ is a flow domain such that M is a positive monoid and E is a nilpotent set of endomorphisms, then this flow domain has unique flows.

Effectively Acyclic Flow Graphs. There are some flow domains that compute flows useful in practice, but which do not guarantee either existence or uniqueness of fixpoints *a priori* for all graphs. For example, the path-counting flow from Example 1 is one where for certain graphs, there exist no solutions to the flow equation (see Figure 5(a)), and for others, there can exist more than one (in Figure 5(b), the nodes marked with x can have any path count, as long as they both have the same value).

In such cases, we explore how to restrict the class of *graphs* we use in our flow-based proofs such that each graph has a unique fixpoint; the difficulty is that this restriction must be respected for composition of our graphs. Here, we study the class of flow domains $(M, +, 0, E)$ such that M is a positive monoid and E is a set of *reduced* endomorphisms (defined below). In such domains we can decompose the flow computations into the various paths in the graph, and achieve unique fixpoints by restricting the kinds of cycles graphs can have.

Definition 11. A flow graph $H = (N, e, fl)$ is effectively acyclic (EA) if for every $1 \leq k$ and $n_1, \dots, n_k \in N$,

$$fl(n_1) \triangleright e(n_1, n_2) \cdots e(n_{k-1}, n_k) \triangleright e(n_k, n_1) = 0.$$

The simplest example of an effectively acyclic graph is one where the edges with non-zero edge functions form an acyclic graph. However, our semantic condition is weaker: for example, when reasoning about two overlaid acyclic lists whose union happens to form a cycle, a product of two path-counting domains will satisfy effective acyclicity because the composition of different types of edges results in the zero function.

Lemma 5. Let $(M, +, 0, E)$ be a flow domain such that M is a positive monoid and E is a closed set of endomorphisms. Given a graph (N, e) over this flow domain and inflow $in: N \rightarrow M$, if there exists a flow graph $H = (N, e, fl)$ that is effectively acyclic, then fl is unique.

While the restriction to effectively acyclic flow graphs guarantees us that the flow is the unique fixpoint of the flow equation, it is not easy to show that modifications to the graph preserve EA while reasoning locally. Even modifying a subgraph to another with the same flow interface (which we know guarantees that it will compose with any context) can inadvertently create a cycle in the larger composite graph. For instance, consider Figure 5(c), that shows a modification to nodes $\{n_3, n_4\}$ (the boxed blue region). The interface of this region is $(\{n_3 \mapsto 1, n_4 \mapsto 1\}, \{n_5 \mapsto 1, n_2 \mapsto 1\})$, and so swapping the edges of n_3 and n_4 preserves this interface. However, the resulting graph, despite composing with the context to form a valid flow graph, is not EA (in this case, it has multiple solutions to the flow equation). This shows that flow interfaces are not powerful enough to preserve effective acyclicity. For a special class of endomorphisms, we show that a local property of the modified subgraph can be checked, which implies that the modified composite graph continues to be EA.

Definition 12. A closed set of endomorphisms $E \subseteq \text{End}(M)$ is called reduced if $e \circ e \equiv \lambda_0$ implies $e \equiv \lambda_0$ for every $e \in E$.

Note that if E is reduced, then no $e \in E$ can be nilpotent. In that sense, this class of instantiations is complementary to the nilpotent class.

Example 6. Examples of flow domains that fall into this class include positive semirings of reduced rings (with the additive monoid of the semiring being the aggregation monoid of the flow domain and E being any set of functions that multiply their argument with a constant flow value). Note that any direct product of integral rings is a reduced ring. Hence, products of the path counting flow domain are a special case.

For reduced endomorphisms, it suffices to check that a modification preserves the flow routed between every pair of source and sink node in order to ensure that it does not create any new cycles in any composite graph.¹⁴

¹⁴ The monoid ordering used in the following definition exists because we are working with a positive monoid.

Definition 13. A flow graph H' is a subflow-preserving extension of H , for which we write $H \lesssim_s H'$, if the following conditions all hold:

- (1) $\text{int}(H) \lesssim \text{int}(H')$
- (2) $\forall n \in H, n' \notin H', m. m \leq \text{inf}(H)(n) \Rightarrow m \triangleright \text{cap}(H)(n, n') = m \triangleright \text{cap}(H')(n, n')$
- (3) $\forall n \in H' \setminus H, n' \notin H', m. m \leq \text{inf}(H')(n) \Rightarrow m \triangleright \text{cap}(H')(n, n') = 0$

This pairwise check, apart from requiring the interface of the modified region to be unchanged, also permits allocating new nodes as long as no flow is routed via the new nodes (condition (3)). We now show that it is sufficient to check that a modification is a subflow-preserving extension to guarantee composition back to an effectively-acyclic composite graph:

Theorem 3. Let $(M, +, 0, E)$ be a flow domain such that M is a positive monoid and E is a reduced set of endomorphisms. If $H = H_1 \odot H_2$ and $H_1 \lesssim_s H'_1$ are all effectively acyclic flow graphs such that $H'_1 \cap H_2 = \emptyset$ and $\forall n \in H'_1 \setminus H_1. \text{outf}(H_2)(n) = 0$, then there exists an effectively acyclic flow graph $H' = H'_1 \odot H_2$ such that $H \lesssim_s H'$.

We define effectively acyclic versions of our flow graph predicates, $N_a(x, H)$ and $\text{Gr}_a(X, H)$, that additionally constrain H to be effectively acyclic. The above theorem yields the following variant of the (REPL) rule for EA graphs:

$$\begin{aligned} & \text{Gr}_a(X'_1, H'_1) * \text{Gr}_a(X_2, H_2) * H = H_1 \odot H_2 * H_1 \lesssim_s H'_1 \\ \models & \exists H'. \text{Gr}_a(X'_1 \uplus X_2, H') * H' = H'_1 \odot H_2 * H \lesssim_s H' \end{aligned} \quad (\text{REPLEA})$$

4.5 Proof of the Harris List

We use the techniques seen in this section in the proof of the Harris list. As the data structure consists of two potentially overlapping lists, we use Lemma 3 to construct a product flow domain of two path-counting flows: one tracks the path count from the head of the main list, and one from the head of the free list. We also work under the effectively acyclic restriction (i.e. we use the N_a and Gr_a predicates), both in order to obtain the desired interpretation of the flow as well as to ensure existence of flows in this flow domain.

We instantiate the framework using the following definitions of parameters:

$$\begin{aligned} fs & := \{\text{key}: k, \text{next}: y, \text{fnext}: z\} \\ \text{edge}(x, fs, v) & := (v = \text{null} ? \lambda_0 : (v = y \wedge y \neq z ? \lambda_{(1,0)} \\ & \quad : (v \neq y \wedge y = z ? \lambda_{(0,1)} : (v = y \wedge y = z ? \lambda_{\text{id}} : \lambda_0)))) \\ \gamma(x, fs, I) & := (I.\text{in}(x) \in \{(1,0), (0,1), (1,1)\}) * (I.\text{in}(x) \neq (1,0) \Rightarrow M(y)) \\ & \quad * (x = ft \Rightarrow I.\text{in}(x) = (_, 1)) * (\neg M(y) \Rightarrow z = \text{null}) \\ \varphi(I) & := I = (\lambda_0[mh \mapsto (1,0)][fh \mapsto (0,1)], \lambda_0) \end{aligned}$$

Here, edge encodes the edge functions needed to compute the product of two path counting flows, the first component tracks path-counts from mh on next edges and the

second tracks path-counts from fh on f_{next} edges¹⁵. The node-local invariant γ says: the flow is one of $\{(1, 0), (0, 1), (1, 1)\}$ (meaning that the node is on one of the two lists, invariant (a)); if the flow is not $(1, 0)$ (the node is not only on the main list, i.e. it is on the free list) then the node is marked (indicated by $M(y)$, invariant (c)); and if the node is ft then it must be on the free list (invariant (d)). The constraint on the global interface, φ , says that the inflow picks out mh and fh as the roots of the lists, and there is no outgoing flow (thus, all non-null edges must stay within the graph, invariant (b)).

Since the Harris list is a concurrent algorithm, we perform the proof in rely-guarantee separation logic (RGSep) [45]. Like in §3, we do not need to modify the semantics of RGSep in any way; our flow-based predicates can be defined and reasoning using our lemmas can be performed in the logic out-of-the-box. For space reasons, the full proof can be found in [Appendix D](#).

5 Related Work

As mentioned in §1, the most closely related work is the flow framework developed by some of the authors in [25]. We here present a simplified and generalized meta theory of flows that makes the approach much more broadly applicable. There were a number of limitations of the prior framework that prevented its application to more general classes of examples.

First, [25] required flow domains to form a semiring; the analogue of edge functions are restricted to multiplication with a constant which must come from the same flow value set. This restriction made it complex to encode many graph properties of interest. For example, one could not easily encode the PIP flow, or a simple flow that counts the number of incoming edges to each node. Our foundational flow framework decouples the algebraic structure defining how flow is *aggregated* from the algebraic structure of the edge functions. In this way, we obtain a more general framework that applies to many more examples, and with simpler flow domains.

Second, in [25], a flow graph did not uniquely determine its inflow (cf. Lemma 1). Correspondingly, [25]’s notion of interface included an *equivalence class* of inflows (all those that induce the same flow values). Since, in [25], the interface also determines which modifications are permitted by the framework, [25] could only handle modifications that preserve the inflow equivalence class. For example, this prevents one from reasoning locally about the removal of a single edge from a graph in certain cases (in particular, like `release` does in the PIP). Our foundational flow framework solves this problem by requiring that the aggregation operation on flow values is cancellative, guaranteeing unique inflows.

To explain this issue in more detail, let us consider the semiring of sets of integers where addition is set union and multiplication is intersection. Further, consider a graph G consisting of nodes $\{n_1, n_2\}$ with a single edge (n_1, n_2) labeled by the set $\{1\}$ and assume that $fl(n_1) = \{1\}$ and $fl(n_2) = \{1, 2\}$. The inflow equivalence class of this graph consists of the two inflow functions $in_1 = \{n_1 \mapsto \{1\}, n_2 \mapsto \{2\}\}$ and $in_2 = \{n_2 \mapsto \{1\}, n_2 \mapsto \{1, 2\}\}$. However, after removing the edge (n_1, n_2) , the two

¹⁵ We use the shorthands $\lambda_{(1,0)} := (\lambda(m_1, m_2). (m_1, 0))$ and $\lambda_{(0,1)} := (\lambda(m_1, m_2). (0, m_2))$, and denote an anonymous existentially-quantified variable by $_$.

inflows are no longer equivalent as they induce different flows. Thus, the new graph has a different set of inflows, and hence a different interface. If the actual inflow provided by G 's context is in_2 , then the removal of the edge would preserve the flow within G . However, the framework in [25] does not allow us to reason locally about this case. In our new framework, requiring that the flow domain is cancellative essentially forces the use of multisets instead of sets in the flow domain (and provides a simpler notion of interface, without equivalence classes). Viewed as inflows over multisets, in_1 and in_2 yield different flows on G and, hence, they belong to two different interfaces. The interface with inflow in_2 permits the removal of the edge whereas the interface with inflow in_1 does not. Requiring a cancellative flow domain (i.e. moving from sets to multisets) solves this problem.

Cancellativity is fundamentally incompatible with [25], which requires the flow domain to form an ω -CPO in order to guarantee the existence of unique flows. For example, in a graph with two nodes n and n' with identity edges between them and all other edges zero (in [25], edges labelled with 1 and 0), if we have $in(n) = 0$ and $in(n') = m$ for some non-zero m , a solution to the flow equation must satisfy $fl(n) = m + fl(n')$. [25] forces such solutions to exist, ruling out cancellativity. To solve this problem, we present a new theory which can optionally guarantee unique flows when desired and show that requiring cancellativity does not limit expressivity.

Next, the proofs of programs shown in [25] depend on a bespoke program logic. This logic requires new reasoning primitives that are not supported by the logics implemented in existing SL-based verification tools. Our general proof technique eliminates the need for a dedicated program logic and can be implemented on top of standard separation logics and existing SL-based tools. Finally, the underlying separation algebra of the original framework makes it hard to use equational reasoning, which is a critical prerequisite for enabling proof automation.

An abundance of SL variants provide complementary mechanisms for modular reasoning about programs (e.g. [21, 40, 42]). Most are parameterized by the underlying separation algebra; our flow-based reasoning technique easily integrates with these existing logics.

Recursive data structures are classically handled in SL using *recursive predicates* [34, 41]. There is a rich line of work in automating such reasoning within decidable fragments (e.g. [2, 12, 18, 23, 36, 38]). However, recursive definitions are problematic for handling e.g. graphs with cycles, sharing and unbounded indegree, overlaid structures and unconstrained traversals.

The most common approach to reason about irregular graph structures in SL is to use iterated separating conjunction [32, 48] and describe the graph as a set of nodes each of which satisfies some local invariant. This approach has the advantage of being able to naturally describe general graphs. However, it is hard to express non-local properties that involve some form of fixpoint computation over the graph structure. One approach is to abstract the program state as a mathematical graph using iterated separating conjunction and then express non-local invariants in terms of the abstract graph rather than the underlying program state [16, 39, 42]. However, a proof that a modification to the state maintains a global invariant of the abstract graph must then often revert back to non-local and manual reasoning, involving complex inductive arguments about paths, transitive

closure, and so on. Our technique also exploits iterated separating conjunction for the underlying heap ownership, with the key benefit that flow interfaces exactly capture the necessary conditions on a modified subgraph in order to compose with *any* context and preserve desired non-local invariants.

In recent work, Wang et al. present a Coq-mechanised proof of graph algorithms in C, based on a substantial library of graph-related lemmas, both for mathematical and heap-based graphs [46]. They prove rich functional properties, integrated with the VST tool. In contrast to our work, a substantial suite of lemmas and background properties are necessary, since these specialise to particular properties such as reachability. We believe that our foundational flow framework could be used to simplify framing lemmas in a way which remains parameteric with the property in question.

Proofs of a number of graph algorithms have been mechanized in various verification tools and proof assistants, including Tarjan’s SCC algorithm [9], union-find [8], Kruskal’s minimum spanning tree algorithm [14], and network flow algorithms [27]. These proofs generally involve non-local reasoning arguments about mathematical graphs.

An alternative approach to using SL-style reasoning is to commit to global reasoning but remain within decidable logics to enable automation [19, 24, 26, 30, 47]. However, such logics are restricted to certain classes of graphs and certain types of properties. For instance, reasoning about reachability in unbounded graphs with two successors per node is undecidable [17]. Recent work by Ter-Gabrielyan et al. [44] shows how to deal with modular framing of *pairwise reachability* specifications in an imperative setting. Their framing notion has parallels to our notion of interface composition, but allows subgraphs to *change* the paths visible to their context. The work is specific to a reachability relation, and cannot express the rich variety of custom graph properties available in our technique.

Dynamic frames [22] (e.g. implemented in Dafny [28]), can be used to explicitly reason about framing of heap information in a first-order logic. However, by itself, this theory does not enable modular reasoning about global graph properties. We believe that the flow framework could in principle be adapted to the dynamic frames setting.

6 Conclusions and Future Work

We have presented the foundational flow framework, enabling local modular reasoning about recursively-defined properties over general graphs. The core reasoning technique has been designed to make minimal mathematical requirements, providing great flexibility in terms of potential instantiations and applications. We identified key classes of these instantiations for which we can provide existence and uniqueness guarantees for the fixpoint properties our technique addresses and demonstrate our proof technique on several challenging examples. As future work, we plan to automate flow-based proofs in our new framework using existing tools that support SL-style reasoning such as Viper [31] and GRASShopper [37].

Acknowledgments. This work is funded in parts by the National Science Foundation under grants CCF-1618059 and CCF-1815633.

References

1. Appel, A.W.: Verified software toolchain. In: NASA Formal Methods. Lecture Notes in Computer Science, vol. 7226, p. 2. Springer (2012)
2. Berdine, J., Calcagno, C., O’Hearn, P.W.: A decidable fragment of separation logic. In: FSTTCS. Lecture Notes in Computer Science, vol. 3328, pp. 97–109. Springer (2004)
3. Bornat, R., Calcagno, C., Yang, H.: Variables as resource in separation logic. *Electron. Notes Theor. Comput. Sci.* **155**, 247–276 (May 2006). <https://doi.org/10.1016/j.entcs.2005.11.059>, <http://dx.doi.org/10.1016/j.entcs.2005.11.059>
4. Brookes, S., O’Hearn, P.W.: Concurrent separation logic. *SIGLOG News* **3**(3), 47–65 (2016)
5. Calcagno, C., Distefano, D., Dubreil, J., Gabi, D., Hooimeijer, P., Luca, M., O’Hearn, P.W., Papakonstantinou, I., Purbrick, J., Rodriguez, D.: Moving fast with software verification. In: NFM. Lecture Notes in Computer Science, vol. 9058, pp. 3–11. Springer (2015)
6. Calcagno, C., O’Hearn, P.W., Yang, H.: Local action and abstract separation logic. In: LICS. pp. 366–378. IEEE Computer Society (2007)
7. Cao, Q., Cuellar, S., Appel, A.W.: Bringing order to the separation logic jungle. In: APLAS. Lecture Notes in Computer Science, vol. 10695, pp. 190–211. Springer (2017)
8. Charguéraud, A., Pottier, F.: Verifying the correctness and amortized complexity of a union-find implementation in separation logic with time credits. *J. Autom. Reasoning* **62**(3), 331–365 (2019)
9. Chen, R., Cohen, C., Lévy, J., Merz, S., Théry, L.: Formal proofs of tarjan’s strongly connected components algorithm in why3, coq and isabelle. In: ITP. LIPIcs, vol. 141, pp. 13:1–13:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2019)
10. Dockins, R., Hobor, A., Appel, A.W.: A fresh look at separation algebras and share accounting. In: APLAS. Lecture Notes in Computer Science, vol. 5904, pp. 161–177. Springer (2009)
11. Dodds, M., Jagannathan, S., Parkinson, M.J., Svendsen, K., Birkedal, L.: Verifying custom synchronization constructs using higher-order separation logic. *ACM Trans. Program. Lang. Syst.* **38**(2), 4:1–4:72 (2016)
12. Enea, C., Lengál, O., Sighireanu, M., Vojnar, T.: SPEN: A solver for separation logic. In: NFM. Lecture Notes in Computer Science, vol. 10227, pp. 302–309 (2017)
13. Harris, T.L.: A pragmatic implementation of non-blocking linked-lists. In: DISC. Lecture Notes in Computer Science, vol. 2180, pp. 300–314. Springer (2001)
14. Haslbeck, M.P.L., Lammich, P., Biendarra, J.: Kruskal’s algorithm for minimum spanning forest. *Archive of Formal Proofs* **2019** (2019)
15. Hoare, C.A.R.: An axiomatic basis for computer programming. *Commun. ACM* **12**(10), 576–580 (1969)
16. Hobor, A., Villard, J.: The ramifications of sharing in data structures. In: POPL. pp. 523–536. ACM (2013)
17. Immerman, N., Rabinovich, A.M., Reps, T.W., Sagiv, S., Yorsh, G.: The boundary between decidability and undecidability for transitive-closure logics. In: CSL. Lecture Notes in Computer Science, vol. 3210, pp. 160–174. Springer (2004)
18. Iosif, R., Rogalewicz, A., Vojnar, T.: Deciding entailments in inductive separation logic with tree automata. In: ATVA. Lecture Notes in Computer Science, vol. 8837, pp. 201–218. Springer (2014)
19. Itzhaky, S., Banerjee, A., Immerman, N., Nanevski, A., Sagiv, M.: Effectively-propositional reasoning about reachability in linked data structures. In: CAV. Lecture Notes in Computer Science, vol. 8044, pp. 756–772. Springer (2013)
20. Jacobs, B., Smans, J., Philippaerts, P., Vogels, F., Penninckx, W., Piessens, F.: Verifast: A powerful, sound, predictable, fast verifier for C and java. In: NASA Formal Methods. Lecture Notes in Computer Science, vol. 6617, pp. 41–55. Springer (2011)

21. Jung, R., Krebbers, R., Jourdan, J., Bizjak, A., Birkedal, L., Dreyer, D.: Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *J. Funct. Program.* **28**, e20 (2018)
22. Kassios, I.T.: Dynamic frames: Support for framing, dependencies and sharing without restrictions. In: FM. *Lecture Notes in Computer Science*, vol. 4085, pp. 268–283. Springer (2006)
23. Katelaan, J., Matheja, C., Zuleger, F.: Effective entailment checking for separation logic with inductive definitions. In: TACAS (2). *Lecture Notes in Computer Science*, vol. 11428, pp. 319–336. Springer (2019)
24. Klarlund, N., Schwartzbach, M.I.: Graph types. In: POPL. pp. 196–205. ACM Press (1993)
25. Krishna, S., Shasha, D.E., Wies, T.: Go with the flow: compositional abstractions for concurrent data structures. *PACMPL* **2**(POPL), 37:1–37:31 (2018)
26. Lahiri, S.K., Qadeer, S.: Back to the future: revisiting precise program verification using SMT solvers. In: POPL. pp. 171–182. ACM (2008)
27. Lammich, P., Sefidgar, S.R.: Formalizing network flow algorithms: A refinement approach in Isabelle/HOL. *J. Autom. Reasoning* **62**(2), 261–280 (2019)
28. Leino, K.R.M.: Dafny: An automatic program verifier for functional correctness. In: LPAR (Dakar). *Lecture Notes in Computer Science*, vol. 6355, pp. 348–370. Springer (2010)
29. Leino, K.R.M., Moskal, M.: Vacid-0: Verification of ample correctness of invariants of data-structures, edition 0. Microsoft Research Technical Report (2010)
30. Madhusudan, P., Qiu, X., Stefanescu, A.: Recursive proofs for inductive tree data-structures. In: POPL. pp. 123–136. ACM (2012)
31. Müller, P., Schwerhoff, M., Summers, A.J.: Viper: A verification infrastructure for permission-based reasoning. In: Jobstmann, B., Leino, K.R.M. (eds.) *Verification, Model Checking, and Abstract Interpretation (VMCAI)*. LNCS, vol. 9583, pp. 41–62. Springer-Verlag (2016)
32. Müller, P., Schwerhoff, M., Summers, A.J.: Automatic verification of iterated separating conjunctions using symbolic execution. In: CAV (1). *Lecture Notes in Computer Science*, vol. 9779, pp. 405–425. Springer (2016)
33. O’Hearn, P.W., Pym, D.J.: The logic of bunched implications. *Bulletin of Symbolic Logic* **5**(2), 215–244 (1999)
34. O’Hearn, P.W., Reynolds, J.C., Yang, H.: Local reasoning about programs that alter data structures. In: CSL. *Lecture Notes in Computer Science*, vol. 2142, pp. 1–19. Springer (2001)
35. Parkinson, M.J., Bierman, G.M.: Separation logic and abstraction. In: Palsberg, J., Abadi, M. (eds.) *Principles of Programming Languages (POPL)*. pp. 247–258. ACM (2005)
36. Piskac, R., Wies, T., Zufferey, D.: Automating separation logic using SMT. In: CAV. *Lecture Notes in Computer Science*, vol. 8044, pp. 773–789. Springer (2013)
37. Piskac, R., Wies, T., Zufferey, D.: Grasshopper - complete heap verification with mixed specifications. In: TACAS. *Lecture Notes in Computer Science*, vol. 8413, pp. 124–139. Springer (2014)
38. Qiu, X., Wang, Y.: A decidable logic for tree data-structures with measurements. In: VMCAI. *Lecture Notes in Computer Science*, vol. 11388, pp. 318–341. Springer (2019)
39. Raad, A., Hobor, A., Villard, J., Gardner, P.: Verifying concurrent graph algorithms. In: APLAS. *Lecture Notes in Computer Science*, vol. 10017, pp. 314–334 (2016)
40. Raad, A., Villard, J., Gardner, P.: Colosl: Concurrent local subjective logic. In: ESOP. *Lecture Notes in Computer Science*, vol. 9032, pp. 710–735. Springer (2015)
41. Reynolds, J.C.: Separation logic: A logic for shared mutable data structures. In: LICS. pp. 55–74. IEEE Computer Society (2002)
42. Sergey, I., Nanevski, A., Banerjee, A.: Mechanized verification of fine-grained concurrent programs. In: PLDI. pp. 77–87. ACM (2015)
43. Sha, L., Rajkumar, R., Lehoczy, J.P.: Priority inheritance protocols: An approach to real-time synchronization. *IEEE Trans. Computers* **39**(9), 1175–1185 (1990)

44. Ter-Gabrielyan, A., Summers, A.J., Müller, P.: Modular verification of heap reachability properties in separation logic. *PACMPL* **3**(OOPSLA), 121:1–121:28 (2019)
45. Vafeiadis, V.: Modular fine-grained concurrency verification. Ph.D. thesis, University of Cambridge, UK (2008)
46. Wang, S., Cao, Q., Mohan, A., Hobor, A.: Certifying graph-manipulating C programs via localizations within data structures. *PACMPL* **3**(OOPSLA), 171:1–171:30 (2019)
47. Wies, T., Muñoz, M., Kuncak, V.: An efficient decision procedure for imperative tree data structures. In: *CADE. Lecture Notes in Computer Science*, vol. 6803, pp. 476–491. Springer (2011)
48. Yang, H.: An example of local reasoning in BI pointer logic: the Schorr-Waite graph marking algorithm. In: *Proceedings of the SPACE Workshop* (2001)
49. Yang, H.: Local reasoning for stateful programs. University of Illinois at Urbana-Champaign (2001)

A Separation Logic

Separation logic (SL), is an extension of Hoare logic [15] that is tailored to perform modular reasoning about programs that manipulate mutable resources. The primary application of SL has been to verify heap-based data structures, but the core of SL is an abstract separation logic (based on the logic of bunched implications (BI) [33]) that can be instantiated to obtain various existing forms of SL by choosing an appropriate resource: any separation algebra (see §2.1).

Heaps. Our separation logic uses standard partial heaps as its semantic model. Let us assume we have the following fixed countably infinite sets: Val , consisting of program values; Loc , consisting of memory addresses; and Field , consisting of field names. Partial heaps are partial maps from addresses to partial maps from field-names to values:

$$\text{Heap} := \{h \mid h: \text{Loc} \rightarrow (\text{Field} \rightarrow \text{Val})\}$$

It is easy to see that, under the disjoint union operator \uplus , and using the empty heap h_\emptyset , $(\text{Heap}, \uplus, h_\emptyset)$ forms a separation algebra.

Programming Language. We consider the following simple imperative programming language:

$C \in \text{Com} ::=$	skip	Empty command
	c	Basic command
	$C_1; C_2$	Sequential composition
	$C_1 + C_2$	Non-deterministic choice
	C^*	Looping
$c ::=$		
	assume (B)	Assume condition
	$x := e$	Variable assignment
	$x := e.f$	Heap dereference
	$e_1.f := e_2$	Heap write
	$x := \mathbf{alloc}()$	Allocate heap cell
	...	

Here, C stands for commands, c for basic commands, x for program variables, e for heap-independent expressions, $f \in \text{Field}$ for field names, and B for boolean expressions. Since we are only concerned with partial correctness in this paper, we can define the more familiar program constructs as the following syntactic shorthands:

$$\begin{aligned} \text{if}(B) C_1 \text{ else } C_2 &::= (\mathbf{assume}(B); C_1) + (\mathbf{assume}(\neg B); C_2) \\ \text{while}(B) C &::= (\mathbf{assume}(B); C)^*; \mathbf{assume}(\neg B) \end{aligned}$$

Assertions. We assume that we start from a standard first-order logic over a signature that includes a countably infinite number of uninterpreted functions and predicates. The only requirement on the underlying logic is that it supports additional uninterpreted sorts, functions and predicates, which can be axiomatised in the pure part of the logic.¹⁶

¹⁶ We will use this power to express all the values associated with flows and flow interfaces.

Let Var be an infinite set of variables (we omit sorts and type-checking from the presentation, for simplicity). The syntax of assertions ϕ is given by the following:

$$\begin{aligned} \phi ::= & P \mid \text{true} \mid \phi \wedge \phi \mid \phi \Rightarrow \phi \mid \exists x. \phi \\ & \mid e \mapsto \{f_1 : e_1, \dots\} \mid \phi * \phi \mid \bigstar_{x \in X} \phi \end{aligned}$$

Here, the first line consists of first order assertions P (called *pure* assertions in the SL world), the always valid assertion *true*, standard boolean connectives, and existential quantification. We can define the remaining boolean connectives and universal quantification as shorthands for the appropriate combination of these. The second line contains the new predicates and connectives introduced by SL (so-called *spatial* assertions). The *points-to* assertion $e \mapsto \{f_1 : e_1, \dots\}$ is a primitive assertion that denotes a heap cell at address e containing fields f_1 with value e_1 , etc. The key feature of SL is the new connective $*$, or *separating conjunction*, that is used to conjoin two *disjoint* parts of the heap. We use the $\bigstar_{x \in X} \phi$ syntax to represent iterated separating conjunction (the bound variable x ranges over a set X).¹⁷

The semantics of the separation logic assertions are defined with respect to an interpretation of (logical and program) variables $i : \text{Var} \rightarrow \text{Val}$. We write $\llbracket e \rrbracket_i$ for the denotation of expression e under interpretation i . In particular, we have:

$$\begin{aligned} h, i \models e \mapsto \{f_1 : e_1, \dots, f_k : e_k\} & \iff h(\llbracket e \rrbracket_i) = \{f_1 \mapsto e_1, \dots, f_k \mapsto e_k\} \\ h, i \models \phi_1 * \phi_2 & \iff \exists h_1, h_2. (h = h_1 \uplus h_2) \\ & \quad \wedge (h_1, i \models \phi_1) \wedge (h_2, i \models \phi_2) \end{aligned}$$

Note that the logic presented here is *garbage-collected* [7] (also known as *intuitionistic*). Thus, the semantics of the points-to assertion $x \mapsto \{f_1 : e_1, \dots, f_k : e_k\}$ does not restrict the heap h to only contain the address x , it only requires x to be included in its domain. This restriction is not essential but simplifies presentation.

Operational Semantics. We give a small-step operational semantics for our programming language. Configurations are either **fault** or a pair (C, σ) of a command C and a state σ (i.e. a heap-interpretation pair). The following rules define a reduction relation \rightarrow between configurations:

$$\begin{array}{c} \text{SEQ1} \\ (\mathbf{skip}; C_2), \sigma \rightarrow C_2, \sigma \end{array} \quad \begin{array}{c} \text{SEQ2} \\ \frac{C_1, \sigma \rightarrow C'_1, \sigma'}{(C_1; C_2), \sigma \rightarrow (C'_1; C_2), \sigma'} \end{array} \quad \begin{array}{c} \text{ASS} \\ \frac{\llbracket B \rrbracket_\sigma}{\mathbf{assume}(B), \sigma \rightarrow \mathbf{skip}, \sigma} \end{array}$$

$$\begin{array}{c} \text{CHO1} \\ (C_1 + C_2), \sigma \rightarrow C_1, \sigma \end{array} \quad \begin{array}{c} \text{CHO2} \\ (C_1 + C_2), \sigma \rightarrow C_2, \sigma \end{array} \quad \begin{array}{c} \text{LOOP} \\ C^*, \sigma \rightarrow (\mathbf{skip} + (C; C^*)), \sigma \end{array}$$

¹⁷ Most presentations of SL also include the *separating implication* connective $-*$. However, logics including $-*$ are harder to automate and usually undecidable. By omitting $-*$ we emphasize that we do not require it to perform flow-based reasoning

While we can also give similar small-step semantics to basic commands (for instance, see [41]), it is easier to understand their axiomatic semantics, presented in the next paragraph.

Soundness of separation logic, especially the frame rule below, relies on the following locality property of the semantics of the programming language. By defining our basic commands via an axiomatic semantics, they automatically satisfy this property, and by construction all composite commands will have the locality property.

Definition 14 (Locality).

- (L1) If $(C, \sigma_1 \odot \sigma) \rightarrow^* \mathbf{fault}$, then $(C, \sigma_1) \rightarrow^* \mathbf{fault}$.
(L2) If $(C, \sigma_1 \odot \sigma) \rightarrow^* (\mathbf{skip}, \sigma_2)$, then either there exists σ'_2 such that $(C, \sigma_1) \rightarrow^* (\mathbf{skip}, \sigma'_2)$ and $\sigma_2 = \sigma \odot \sigma'_2$, or $(C, \sigma_1) \rightarrow^* \mathbf{fault}$.

Proof Rules. As with Hoare logic, programs are specified in separation logic by Hoare triples.

Definition 15 (Hoare Triple). We say $\models \{\phi\} C \{\psi\}$ if for every state σ such that $\sigma \models \phi$ we have (1) $(C, \sigma) \not\rightarrow^* \mathbf{fault}$, and (2) for every state σ' such that $(C, \sigma) \rightarrow^* (\mathbf{skip}, \sigma')$, $\sigma' \models \psi$.

In the above definition, \rightarrow^* is the reflexive transitive closure of the reduction relation \rightarrow . Intuitively, the judgment $\{\phi\} C \{\psi\}$ means that if a command C is executed on a state satisfying the precondition ϕ , then it executes without faults. Moreover, if C terminates, then the resulting state satisfies the postcondition ψ (thus, this is a *partial correctness criterion*).

Separation logic inherits the standard Floyd-Hoare structural proof rules, and the rule of consequence:

$$\begin{array}{c}
\text{SL-SKIP} \\
\frac{}{\vdash \{\phi\} \mathbf{skip} \{\phi\}} \\
\\
\text{SL-CHOICE} \\
\frac{\vdash \{\phi\} C_1 \{\psi\} \quad \vdash \{\phi\} C_2 \{\psi\}}{\vdash \{\phi\} C_1 + C_2 \{\psi\}} \\
\\
\text{SL-CONSEQ} \\
\frac{P' \Rightarrow \phi \quad \vdash \{\phi\} C \{\psi\} \quad \psi \Rightarrow Q'}{\vdash \{P'\} C \{Q'\}} \\
\\
\text{SL-SEQ} \\
\frac{\vdash \{\phi\} C_1 \{\psi\} \quad \vdash \{\psi\} C_2 \{\rho\}}{\vdash \{\phi\} C_1; C_2 \{\rho\}} \\
\\
\text{SL-LOOP} \\
\frac{\vdash \{\phi\} C \{\phi\}}{\vdash \{\phi\} C^* \{\phi\}}
\end{array}$$

The scalability of SL-based reasoning arises due to the following *frame rule*:

$$\text{SL-FRAME} \\
\frac{\vdash \{\phi\} C \{\psi\}}{\vdash \{\phi * \rho\} C \{\psi * \rho\}}$$

The frame rule allows one to lift a proof that a command C executes safely on a state satisfying ϕ , producing a state satisfying ψ if it terminates, to the setting where an

additional resource ρ (the *frame*) is present. Since C was safe when given only ϕ , it does not access any resources outside ϕ ; hence, ρ is untouched in the postcondition. The soundness of the frame rule relies on the disjointness of resources enforced by the separating conjunction operator $*$.¹⁸

For the basic commands of the programming language, one can give *small axioms*, proof rules that specify the minimum resource they need in order to execute safely. The effect of basic commands on more complex states can be derived from these and the frame rule. Here are some of the small axioms:

$$\begin{array}{l} \text{SL-ASSIGN} \\ \vdash \{\psi[x \mapsto e]\} x := e \{\psi\} \\ \\ \text{SL-WRITE} \\ \vdash \{e_1 \mapsto \{f: _, \dots\}\} e_1.f := e_2 \{e_1 \mapsto \{f: e_2, \dots\}\} \\ \\ \text{SL-READ} \\ \vdash \{e \mapsto \{f: z, \dots\} * e = y\} x := e.f \{y \mapsto \{f: z, \dots\} * x = z\} \end{array}$$

Note that we write $\psi[x \mapsto e]$ for the assertion ψ where all occurrences of x are replaced with e , and $_$ for an anonymous existential variable (to denote expressions we do not care about).

Together with standard axioms of first-order logic, the proof rules presented above are known to be complete [49].¹⁹ In other words, all valid Hoare triples can be derived by an appropriate combinations of these axioms.

B Expressivity of Flows

We now give a few examples of flows to demonstrate the range of data structures whose properties can be expressed as local constraints on each node's flow.

We start by describing a few interesting examples of flows that capture common graph properties. The path-counting flow defined in §2.2 is a very useful flow for describing the shape of common structures, e.g. lists (singly and doubly linked, cyclic), trees, and (by using product flow constructions) nested and overlaid combinations of these. By considering products with flows for data properties, we can also describe structures such as sorted lists, binary heaps, and search trees.

The next flow is similar to the PIP flow defined in §2.2 and can be used to specify the correctness of algorithms such as Dijkstra's shortest paths algorithm.

Definition 16 (Shortest Path Flow). *The shortest path flow uses the flow domain $(\mathbb{N}^C, \cup, \emptyset, \{\lambda_n \mid n \in C\})$ of multisets over costs $C = \mathbb{N} \cup \{\infty\}$ where $\lambda_n(S) = \{n + \min(S)\}$.*

¹⁸ The frame rule relies on a side condition that the program variables modified by C do not overlap with the free variables in ρ , but this condition can be omitted using the “variables as resource” technique [3].

¹⁹ Note that Yang's completeness result depends crucially on the separating implication \multimap being included in the assertion language.

Given a flow graph $H = (N, e, fl)$ over this domain and a cost labeling function $c : N \times N \rightarrow C$ for edges, if $e(n, n')$ is $\lambda_{c(n, n')}$ and $\text{inf}(H) = (\lambda n. (n = s ? \{0\} : \emptyset))$ for some source node s , then $fl(n)$ is the multiset of costs of all shortest paths from s to n via n 's predecessors. That is, the cost of the shortest path from s to n is $\min(fl(n))$.

The next flow can be used to reason about reachability properties in graphs.

Definition 17 (Inverse Reachability Flow). Consider the flow domain $(\mathbb{N}^{2^{\mathfrak{N}}}, \cup, \emptyset, E)$, consisting of the monoid of multisets of sets of nodes under multiset union and edge functions E containing λ_0 and for every $n \in \mathfrak{N}$ the function

$$\lambda_n(S) := \{P \mapsto (n \in P ? S(P \setminus \{n\}) : 0)\}.$$

Given a flow graph $H = (N, e, fl)$, if $e(n, n') = \lambda_n$ and $\text{inf}(H) = (\lambda n. (n = r ? \{\emptyset\} : \emptyset))$, then $fl(n)$ is a multiset containing, for each simple path in H from r to n , the set of all nodes occurring on that path.

Finally, we demonstrate the full generality of the flow equation in terms of its ability to capture global graph properties. To this end, we define a *universal* flow that computes, at each node, sufficient information to reconstruct the entire graph. This shows that flows are powerful enough to capture any graph property of interest.

Definition 18 (Universal Flow). Say we are given a set of nodes $N \subseteq \mathfrak{N}$ and a function $\epsilon : N \times \mathfrak{N} \rightarrow A$ labelling each pair of nodes from some set A (for instance, to encode an unlabelled graph, $A = \{0, 1\}$ and $\epsilon(n, n')$ is 1 iff an edge is present in the graph). Consider the flow domain $(\mathbb{N}^{2^{\mathfrak{N}} \times \mathfrak{N} \times A}, \cup, \emptyset, E)$, consisting of the monoid of multisets of sets of tuples (n, n', a) of edges (n, n') and labels $a \in A$ under multiset union and edge functions E containing λ_0 and for every $n, n' \in \mathfrak{N}, a \in A$ the function

$$\lambda_{n, n', a}(S) := \{P \mapsto ((n, n', a) \in P ? S(P \setminus \{(n, n', a)\}) : 0)\}.$$

Given a flow graph $H = (N, e, fl)$, if $e(n, n') = \lambda_{n, n', \epsilon(n, n')}$ and $\text{inf}(H) = (\lambda n. \{\emptyset\})$, then $fl(n)$ is a multiset containing, for each simple path in H ending at n , a set of all edge-label tuples of edges occurring on that path.

To see why the universal flow computes the entire graph at each node, let us look at the edge functions in more detail. The way to think of $e(n, n') = \lambda_{n, n', \epsilon(n, n')}$ is that it looks at each path P' in the input multiset S and if P' does not contain the tuple $(n, n', \epsilon(n, n'))$ then it adds the tuple to P' and adds the resulting path P to the output multiset. In order to convert this procedure into a multiset comprehension style definition, the formal definition above starts from each path P in the output multiset and works backward (i.e. $P' = P \setminus \{(n, n', \epsilon(n, n'))\}$).

To understand the flow computation, let us start with the inflow to a node n , the singleton multiset containing the empty set $\{\emptyset\}$, and track its progress through a path. For every n' , the edge function $e(n, n')$ acts on $\{\emptyset\}$ and propagates the singleton multiset $\{(n, n', \epsilon(n, n'))\}$. In this way, if we consider a sequence of (distinct) edges $(n_1, n_2), \dots, (n_{k-1}, n_k)$, then this value becomes

$$\{(n_1, n_2, \epsilon(n_1, n_2)), \dots, (n_{k-1}, n_k, \epsilon(n_{k-1}, n_k))\} .$$

However, the minute we follow an edge (n_i, n_{i+1}) that has occurred on the path before, the edge function $e(n_i, n_{i+1})$ will send this value to the empty multiset \emptyset . Thus, the flow at each node n turns out to be the multiset containing sets of edge-label tuples for each simple path in the graph.²⁰ Note that we label all pairs of nodes in the graph by edges of the form $\lambda_{n,n',\epsilon(n,n')}$. This means that $fl(n)$ will contain one set for every sequence of pairs of nodes in the graph, even those corresponding to edges that do not “exist” in the original graph ϵ . From this information, one can easily reconstruct all of ϵ and hence any graph property of the global graph.

The power of the universal flow to capture any graph property comes with a cost: any modification to the graph entails a global modification to *every* node in the graph. This means that we lose all powers of local reasoning, and revert to expensive global reasoning about the program. This is to be expected, however, because the universal flow captures all details of the graph, even ones that are possibly irrelevant to the correctness of the program at hand. The art of using flows is to carefully define a flow that captures exactly the necessary global information needed to locally prove correctness of a given program.

For example, the inverse reachability flow is a simplified version of the universal flow in that for each edge it only keeps track of the source node. By capturing less information about the global graph, however, this flow permits more modifications: for instance, one can swap the order of two nodes in a simple path and only update the flows of the two nodes modified. This is an example of carefully tuning the flow domain to match the modifications performed by the program.

C The PIP

In order to expose field values in the top level specification (e.g. to say that `acquire` results in the appropriate edge) we extend the signatures of our core predicates to allow extra custom parameters: $\gamma(x, fs, m, \dots)$ and $N(x, H, \dots)$. For the PIP, we instantiate the framework as follows, where η is a function from nodes to nodes storing the values

²⁰ This flow domain has the property that any graph has a unique solution to the flow equation (see the nilpotent class in §4.4).

of the next fields (as enforced by the last line of γ):

$$fs := \{\text{next}: y, \text{curr_prio}: q, \text{def_prio}: q^0, \text{prios}: Q\}$$

$$\text{edge}(x, fs, z) := \begin{cases} (\lambda m. \max(m \cup \{q^0\})) & \text{if } z = y \neq \text{null} \\ \lambda_0 & \text{otherwise} \end{cases}$$

$$\begin{aligned} \gamma(x, fs, m, \eta) &:= q^0 \geq 0 * (\forall q' \in Q. q' \geq 0) \\ &* m = Q * q = \{\max(Q \cup \{q^0\})\} \\ &* \eta(x) = y * \eta(x) \neq x \end{aligned}$$

$$\varphi(I) := I = (\{- \mapsto \emptyset\}, \{- \mapsto \emptyset\})$$

Thus, $\mathbf{N}(x, H, \eta)$ describes a node x abstracted by flow graph H and whose next field is $\eta(x)$.

```

1 // Let  $\delta(m, q_1, q_2) := m \setminus (q_1 \geq 0 ? \{q_1\} : \emptyset) \cup (q_2 \geq 0 ? \{q_2\} : \emptyset)$ 
2 // Let  $\Delta(I, n, q_1, q_2) := (\{n \mapsto \delta(I.in(n), q_1, q_2)\}, I.out)$ 
3
4 method update(n: Ref, from: Int, to: Int)
5   requires  $\mathbf{N}(n, I_n, \eta) * \mathbf{Gr}(X \setminus \{n\}, I', \eta) * I = I'_n \oplus I' * \varphi(I) * n \in X$ 
6   requires  $I'_n = \Delta(I_n, n, \text{from}, \text{to}) * \text{from} \neq \text{to}$ 
7   ensures  $\mathbf{Gr}(X, I, \eta)$ 
8 {
9   n.prios := n.prios \ {from}
10  if (to >= 0) {
11    n.prios := n.prios  $\cup$  {to}
12  }
13  from := n.curr_prio
14  n.curr_prio := max(n.prios  $\cup$  {n.def_prio})
15  to := n.curr_prio
16
17  if (from != to && n.next != null) { // Let  $n' := n.next$ 
18     $\left\{ \begin{array}{l} \exists I'_{n'}, I''_n, I_1. \mathbf{N}(n, I''_n, \eta) * \mathbf{N}(n', I_{n'}, \eta) * \mathbf{Gr}(X \setminus \{n, n'\}, I_1, \eta) * I = I''_n \oplus I'_{n'} \oplus I_1 \\ * I''_n = (I'_n.in, \{n' \mapsto \{to\}\}) * I'_{n'} = \Delta(I_{n'}, n', \text{from}, \text{to}) * \text{from} \neq \text{to} \end{array} \right\}$ 
19     $\models$ 
20     $\left\{ \begin{array}{l} \exists I'_{n'}, I'. \mathbf{N}(n', I_{n'}, \eta) * \mathbf{Gr}(X \setminus \{n'\}, I', \eta) * I' \oplus I'_{n'} = I \\ * I'_{n'} = \Delta(I_{n'}, n', \text{from}, \text{to}) * \text{from} \neq \text{to} \end{array} \right\}$ 
21    update(n.next, from, to)
22  }
23 }
24
25 method acquire(p: Ref, r: Ref)
26   requires  $\mathbf{Gr}(X, I, \eta) * \varphi(I)$ 
27   requires  $p \in X * r \in X * p \neq r * \eta(p) = \text{null}$ 

```

```

28 ensures Gr( $X, I, \eta'$ )
29 ensures ( $\eta(r) = \text{null} ? \eta' = \eta[r \mapsto p] : \eta' = \eta[p \mapsto r]$ )
30 {
31    $\{\exists I_r, I_p, I_1. \mathbf{N}(r, I_r, \eta) * \mathbf{N}(p, I_p, \eta) * \text{Gr}(X \setminus \{r, p\}, I_1, \eta) * I = I_r \oplus I_p \oplus I_1 * \varphi(I)\}$ 
32   if (r.next == null) {
33     // Let  $q_r = \text{r.curr\_prio}$  ( $\geq 0$  due to  $\mathbf{N}$ )
34      $\left\{ \begin{array}{l} \exists I_r, I_p, I_1. \mathbf{N}(r, I_r, \eta) * \mathbf{N}(p, I_p, \eta) * \text{Gr}(X \setminus \{r, p\}, I_1, \eta) * I = I_r \oplus I_p \oplus I_1 \\ * q_r \geq 0 * I_r.out = \lambda_0 * \dots \end{array} \right\}$ 
35     r.next := p
36     // Let  $\eta' = \eta[r \mapsto p]$ 
37      $\left\{ \begin{array}{l} \exists I_r, I'_r, I_p, I_1. \mathbf{N}(r, I'_r, \eta') * \mathbf{N}(p, I_p, \eta') * \text{Gr}(X \setminus \{r, p\}, I_1, \eta) * I = I_r \oplus I_p \oplus I_1 \\ * I'_r = (I_r.in, \{p \mapsto \{q_r\}\}) * q_r \geq 0 * I_r.out = \lambda_0 * \dots \end{array} \right\}$ 
38      $\models$ 
39      $\left\{ \begin{array}{l} \exists I_p, I'_p, I_2. \mathbf{N}(p, I_p, \eta') * \text{Gr}(X \setminus \{p\}, I_2, \eta') * I = I'_p \oplus I_2 \\ * I'_p = (\{p \mapsto \delta(I_p.in(p), -1, q_r)\}, I_p.out) * \dots \end{array} \right\}$ 
40     update(p, -1, r.curr_prio)
41      $\{\text{Gr}(X, I, \eta')\}$ 
42   } else {
43     p.next := r
44     update(r, -1, p.curr_prio)
45   }
46 }
47
48 method release(p: Ref, r: Ref)
49 requires Gr( $X, I, \eta$ ) *  $\varphi(I)$ 
50 requires  $p \in X * r \in X * p \neq r * \eta(r) = p$ 
51 ensures Gr( $X, I, \eta'$ )
52 ensures  $\eta' = \eta[r \mapsto \text{null}]$ 
53 {
54   r.next := null
55   update(p, r.curr_prio, -1)
56 }

```

D The Harris List

We perform the proof of the Harris list in rely-guarantee separation logic (RGSep) [45]. RGSep is parametrised by the program states (any separation algebra), the language of assertions (a variant of separation logic), and the programming language (as long as the basic commands are *local*, see §A). Unlike [25], where RGSep was instantiated with a bespoke separation algebra, assertion language, and programming language with custom constructs for flows, in this paper we instantiate RGSep with the standard partial heap separation algebra and the standard separation logic assertion language from §A. We do, however, need to switch to a concurrent programming language, so we adopt the simple imperative language used in by the original RGSep presentation [45]. As we did in §3, we only need to define flow predicates \mathbf{N} and Gr within the logic and assume the flow lemmas from Figure 2 in order to perform flow-based proofs in RGSep.

For the Harris, we instantiate the framework as follows, where we extend the γ and N predicates to also keep track of the mark status (the M predicate encodes whether a reference is marked), next and fnext values of each node:

$$fs := \{\text{key}: k, \text{next}: y, \text{fnext}: z\}$$

$$\begin{aligned} \text{edge}(x, fs, v) := & (v = \text{null} ? \lambda_0 \\ & : (v = y \wedge y \neq z ? \lambda_{(1,0)} \\ & : (v \neq y \wedge y = z ? \lambda_{(0,1)} \\ & : (v = y \wedge y = z ? \lambda_{\text{id}} : \lambda_0)))) \end{aligned}$$

$$\begin{aligned} \gamma(x, fs, I, m, x_n, x_f) := & (I.\text{in}(x) \in \{(1,0), (0,1), (1,1)\}) \\ & * (I.\text{in}(x) \neq (1,0) \Rightarrow M(y)) \\ & * (x = \text{ft} \Rightarrow I.\text{in}(x) = (_, 1)) \\ & * (\neg M(y) \Rightarrow z = \text{null}) \\ & * (M(y) \Rightarrow m = \blacklozenge_) * (\neg M(y) \Rightarrow m = \blacklozenge) \\ & * x_n = y * x_f = z \end{aligned}$$

$$\begin{aligned} \varphi(I) := & I.\text{in} = \{mh \mapsto (1,0), fh \mapsto (0,1), _ \mapsto (0,0)\} \\ & * I.\text{out} = \{_ \mapsto (0,0)\} \end{aligned}$$

$$\begin{aligned} N(x, I, m, x_n, x_f) := & x \mapsto fs * \gamma(x, fs, I, m, x_n, x_f) \\ & * \text{dom}(I) = \{x\} * \forall y. I.\text{out}(y) = \text{edge}(x, fs, y)(I.\text{in}(x)) \end{aligned}$$

Note that we extend γ with extra parameters m , x_n , and x_f that keep track of the mark status (the M predicate encodes whether a reference is marked), next and fnext values of x . We also expose these parameters in the N predicate. In our proof, we ignore these additional parameters to N when we do not care about them (i.e. $N(x, I) = N(x, I, _, _, _)$, etc.).

D.1 Actions

RGSep consists of two kinds of assertions: boxed assertions $\boxed{\phi}$ talk about the shared state among threads, and unboxed assertions ϕ talk about thread-local state. An RGSep proof requires an intermediate assertion in between every two atomic modifications to the shared state, along with a stability proof that this intermediate assertion is preserved by interference of other threads. Interference is formally specified via *actions*, two-state relations that describe modifications performed by threads to the shared state.

$$\begin{aligned}
\mathsf{N}(l, I_l, \diamond) &\rightsquigarrow \mathsf{N}(l, I'_l, \diamond) * \mathsf{N}(n, I_n) && \text{(Insert)} \\
&\quad * I_l \lesssim (I'_l \oplus I_n) \\
\mathsf{N}(r, I_r, \diamond) &\rightsquigarrow \mathsf{N}(r, I_r, \blacklozenge_t) && \text{(Mark)} \\
\mathsf{N}(ft, I_f, \blacklozenge_-, x, null) * \mathsf{N}(r, I_r, \blacklozenge_t, y, null) &\rightsquigarrow \mathsf{N}(ft, I'_f, \blacklozenge_-, x, r) * \mathsf{N}(r, I'_r, \blacklozenge_t, y, null) \\
&\quad * (I_f \oplus I_r) = (I'_f \oplus I'_r) && \text{(Link)} \\
\mathsf{N}(l, I_l, \diamond, r, -) * \mathsf{N}(r, I_r, \blacklozenge_-, x, y) &\rightsquigarrow \mathsf{N}(l, I'_l, \diamond, x, -) * \mathsf{N}(r, I'_r, \blacklozenge_-, x, y) \\
&\quad * (I_l \oplus I_r) = (I'_l \oplus I'_r) && \text{(Unlink)}
\end{aligned}$$

D.2 Proof

Since we work under the effectively acyclic restriction, we use the predicates N_a and Gr_a in our proof. As we did with N and Gr , we overload N_a and Gr_a to talk about flow interfaces instead of flow graphs:

$$\begin{aligned}
\mathsf{N}_a(x, I) &:= \exists H. \mathsf{N}_a(x, H) * I = \text{int}(H) \\
\mathsf{Gr}_a(X, I) &:= \exists H. \mathsf{Gr}_a(x, H) * I = \text{int}(H)
\end{aligned}$$

The only place where effective acyclicity reasoning comes in is when we reason about modifications to the shared state: we then have to show that the modified flow graph is a subflow-preserving extension (so that we can use the rule (REPLEA)). We show some examples of this below with purple annotations in braces. Finally, the proof also requires showing stability of all intermediate assertions in blue. This is straightforward and follows the proof method used by [25, 45], so we omit these details here.

```

1 procedure search(k: Key) returns (l: Ref, r: Ref)
2 requires  $\exists X, I. \mathsf{Gr}_a(X, I) * \varphi(I)$ 
3 ensures  $\exists X, I. \mathsf{Gr}_a(X, I) * \varphi(I) * \{l, r\} \subseteq X$ 
4 {
5    $\left\{ \exists X, I_{mh}, I_1. \mathsf{N}_a(mh, I_{mh}) * \mathsf{Gr}_a(X \setminus \{mh\}, I_1) * \varphi(I_{mh} \oplus I_1) * \{l, r\} \subseteq X \right\}$ 
6   l := mh; r := mh
7   var n := head.next
8
9   while (isMarkedRef(n) || r.key < k)
10  invariant  $\exists X, I. \mathsf{Gr}_a(X, I) * \varphi(I) * \{l, r\} \subseteq X * (\neg M(n) \Rightarrow n \in X)$ 
11  {
12    if (isMarkedRef(n)) { // r marked
13       $\left\{ \exists X, I_l, I_1. \mathsf{N}_a(l, I_l) * \mathsf{Gr}_a(X \setminus \{l\}, I_1) * \varphi(I_l \oplus I_1) \right\}$ 
14      if (CAS(l.next, r, unmarked(n))) { // try to unlink r
15        // Success: now l.next == n, l still unmarked
16        r := unmarked(n)

```

```

17   } else {
18     // either something inserted in between l and r or marked l, restart
19     search(k)
20   }
21   } else {
22     l := r
23     r := n
24   }
25    $\{\exists X, I_r, I_1. N_a(r, I_r) * Gr_a(X \setminus \{r\}, I_1) * \varphi(I_r \oplus I_1) * l \in X\}$ 
26   n := r.next
27 }
28 return (l, r)
29 }
30
31 procedure insert(k: Key)
32 requires  $\{\exists X, I. Gr_a(X, I) * \varphi(I)\}$ 
33 ensures  $\{\exists X, I. Gr_a(X, I) * \varphi(I)\}$ 
34 {
35   Node n := new Node(k, null, null)
36    $\{\exists X, I. Gr_a(X, I) * \varphi(I) * N_a(n, \_, \diamond, null, null)\}$ 
37
38   while (true)
39     invariant  $\{\exists X, I. Gr_a(X, I) * \varphi(I) * N_a(n, \_, \diamond, null, null)\}$ 
40     {
41       var l, r := search(k)
42        $\{\exists X, I. Gr_a(X, I) * \varphi(I) * \{l, r\} \subseteq X * N_a(n, \_, \diamond, null, null)\}$ 
43       if (r.key == k) {
44         free(n)
45         return false
46       }
47       n.next := r
48        $\{\exists X, I_l, I_r, I_1. N_a(l, I_l) * Gr_a(X \setminus \{l\}, I_1) * \varphi(I_l \oplus I_1) * N_a(n, \_, \diamond, null, null)\}$ 
49       if (CAS(l.next, r, n)) {
50          $\{N_a(l, H'_l) * N_a(n, H'_n) * Gr_a(X \setminus \{l\}, H_1) * H_l \lesssim_s (H'_l \oplus H'_n) * \varphi(\text{int}(H_l \oplus H_1))\}$ 
51         return true
52       }
53     }
54 }
55
56 procedure delete(k: Key)
57 requires  $\{\exists X, I. Gr_a(X, I) * \varphi(I)\}$ 
58 ensures  $\{\exists X, I. Gr_a(X, I) * \varphi(I)\}$ 
59 {
60   var l, r, n

```



```

61
62 while (true)
63   invariant  $\exists X, I. \text{Gr}_a(X, I) * \varphi(I)$ 
64   {
65     l, r := search(k)
66      $\{ \exists X, I. \text{Gr}_a(X, I) * \varphi(I) * \{l, r\} \subseteq X \}$ 
67     n := r.next
68     if (r.key != k) {
69       return false
70     }
71     if (!isMarkedRef(n)) { // r unmarked
72        $\{ \exists X, I_r, I_1. \text{N}_a(r, I_r) * \text{Gr}_a(X \setminus \{r\}, I_1) * \varphi(I_r \oplus I_1) * l \in X * \neg M(n) \}$ 
73       if (CAS(r.next, n, marked(n))) { // mark r
74          $\{ \text{N}_a(r, I_r, \blacklozenge_t, \text{null}) * \text{Gr}_a(X \setminus \{r\}, I_1) * \varphi(I_r \oplus I_1) * \{l, r\} \subseteq X \}$ 
75         break
76       }
77     } // if r already marked, we should have returned false, so retry
78   }
79    $\{ \exists X, I. \text{N}_a(r, I_r, \blacklozenge_t, \text{null}) * \text{Gr}_a(X \setminus \{r\}, I_1) * \varphi(I_r \oplus I_1) * l \in X * r \neq \text{ft} \}$ 
80   // Try to unlink r from main list
81   CAS(l.next, r, n) // If this fails, next search will unlink
82    $\{ \exists X, I. \text{N}_a(r, I_r, \blacklozenge_t, \text{null}) * \text{Gr}_a(X \setminus \{r\}, I_1) * \varphi(I_r \oplus I_1) * l \in X * r \neq \text{ft} \}$ 
83   // Link r to free list
84   while (!CAS(ft.next, null, r)) {}
85    $\{ \exists X, I. \text{N}_a(r, I_r, \blacklozenge_t, \text{null}) * \text{Gr}_a(X \setminus \{r\}, I_1) * \varphi(I_r \oplus I_1) * l \in X * r \neq \text{ft} * I_r.in(r) = (\_, 1) \}$ 
86   ft := r
87
88   return true
89 }

```

E Proofs

This section contains proofs of all lemmas and theorems from the main body of the paper. For clarity, we repeat the lemma and theorem statements before giving the proofs, and used the same numbers as in the body of the paper.

E.1 Section §2

Lemma 1. *Given a flow graph (N, e, fl) , there exists a unique inflow in such that $\text{FlowEqn}(in, e, fl)$.*

Proof. Suppose in and in' are two solutions to $\text{FlowEqn}(_, e, fl)$. Then, for any n ,

$$fl(n) = in(n) + \sum_{n' \in \text{dom}(in)} fl(n') \triangleright e(n', n) = in'(n) + \sum_{n' \in \text{dom}(in')} fl(n') \triangleright e(n', n)$$

which, by cancellativity of the flow domain, implies that $in(n) = in'(n)$. \square

The following lemma justifies the use of the interfaces of a singleton flow graphs for writing flow-based specs, as done in §3. It says that a singleton flow interface containing n captures the flow of n and the outflow values propagated by n 's edges:

Lemma 6. *For any flow graph $H = (N, e, fl)$ and $n, n' \in N$, if $e(n, n) = \lambda_0$ then $int(H_n).in(n) = fl(n)$ and $int(H_n).out(n') = fl(n) \triangleright e(n, n')$.*

Proof. Follows directly from (FlowEqn) and the definition of outflow. \square

Lemma 2. $int(H_1) = I_1 \wedge int(H_2) = I_2 \Rightarrow int(H_1 \odot H_2) = I_1 \oplus I_2$.

Proof. If $H_1 \odot H_2$ is defined and has interface I , then we show that $I_1 \oplus I_2$ is defined and equal to I . Let $H_i = (N_i, e_i, fl_i), I = (in, out), I_1 = (in_1, out_1)$, and $I_2 = (in_2, out_2)$. Since $H = H_1 \odot H_2 \in FG$ and $inf(H) = I.in = in$, we know by definition that $\forall i \neq j \in \{1, 2\}, n \in H_i$,

$$\begin{aligned}
fl(n) &= in(n) + \sum_{n' \in H} fl(n') \triangleright e(n', n) \\
\iff fl_i(n) &= in(n) + \sum_{n' \in H} fl(n') \triangleright e(n', n) \\
\iff in_i(n) + \sum_{n' \in H_i} e_i(n', n, fl_i(n')) & \\
&= in(n) + \sum_{n' \in H_i} e_i(n', n, fl_i(n')) + \sum_{n' \in H_j} e_j(n', n, fl_j(n')) \\
\iff in_i(n) &= in(n) + \sum_{n' \in H_j} e_j(n', n, fl_j(n')) \quad (\text{By cancellativity}) \\
\iff in_i(n) &= in(n) + out_j(n).
\end{aligned}$$

Secondly, let $H = H_1 \odot H_2$ and note that

$$\begin{aligned}
out(n) &:= \sum_{n' \in H} fl(n') \triangleright e(n', n) \\
&= \sum_{n' \in H_1} fl_1(n') \triangleright e_1(n', n) + \sum_{n' \in H_2} fl_2(n') \triangleright e_2(n', n) \\
&= out_1(n) + out_2(n).
\end{aligned}$$

As $H = H_1 \odot H_2$ implies $dom(H_1) \cap dom(H_2) = \emptyset$, this proves that $I_1 \oplus I_2 = I$.

Conversely, if $I_1 \oplus I_2$ is defined and equal to I then we show that $H_1 \odot H_2$ is defined and has interface I . First, $I_1 \cap I_2 = \emptyset$, so we know that the graphs are disjoint. Note that the proof above works in both directions, so

$$\begin{aligned}
\forall i \neq j \in \{1, 2\}, n \in I_i. I_i.in(n) &= I.in(n) + I_j.out(n) \\
\Rightarrow fl(n) &= in(n) + \sum_{n' \in H} fl(n') \triangleright e(n', n).
\end{aligned}$$

This tells us that $H = H_1 \odot H_2 \in \text{FG}$ and $\text{inf}(H) = \text{in}$. From above, we also know that $\text{out}(n) = \text{out}_1(n) + \text{out}_2(n)$, so the interface composition condition $I.\text{out}(n) = I_1.\text{out}(n) + I_2.\text{out}(n)$ gives us $\text{outf}(H) = \text{out}$. \square

Theorem 1. *The flow interface algebra $(\text{FI}, \oplus, I_\emptyset)$ is a separation algebra.*

Proof. We prove commutativity first, as it is used in the proof of associativity:

– \oplus is commutative:

This follows from the symmetry in the definition of \oplus and the commutativity of the flow domain operator $+$.

– \oplus is associative, i.e. $I_1 \oplus (I_2 \oplus I_3) = (I_1 \oplus I_2) \oplus I_3$:

Note that if any two of the interfaces I_1, I_2 , and I_3 are not disjoint, then both sides of the equation are equal to \perp . We now show that if the LHS is defined, then the RHS is also defined and equal to it. Let $I_{23} = I_2 \oplus I_3$, and $I = I_1 \oplus I_{23}$. Define $I_{12} = ((\lambda n. I.\text{in}(n) + I_3.\text{out}(n)), (\lambda n'. I_1.\text{out}(n') + I_2.\text{out}(n)))$. We first show that $I_{12} = I_1 \oplus I_2$. We know that the interfaces are disjoint, so let us check that the inflows are compatible. For $n \in I_1$,

$$\begin{aligned} I_{12}.\text{in}(n) + I_2.\text{out}(n) &= I.\text{in}(n) + I_3.\text{out}(n) + I_2.\text{out}(n) \\ &= I.\text{in}(n) + I_{23}.\text{out}(n) && (\text{As } I_{23} = I_2 \oplus I_3) \\ &= I_1.\text{in}(n). && (\text{As } I = I_1 \oplus I_{23}) \end{aligned}$$

On the other hand, for $n \in I_2$,

$$\begin{aligned} I_{12}.\text{in}(n) + I_1.\text{out}(n) &= I.\text{in}(n) + I_3.\text{out}(n) + I_1.\text{out}(n) \\ &= I_{23}.\text{in}(n) + I_3.\text{out}(n) && (\text{As } I = I_{23} \oplus I_1) \\ &= I_2.\text{in}(n). && (\text{As } I_{23} = I_2 \oplus I_3) \end{aligned}$$

Finally, the condition on the outflow follows by definition of I_{12} .

We now show that $I = I_{12} \oplus I_3$. Again, the disjointness condition is satisfied by assumption, so let us check the inflows. If $n \in I_{12}$, $I.\text{in}(n) + I_3.\text{out}(n) = I_{12}.\text{in}(n)$ by definition of I_{12} . And if $n \in I_3$,

$$\begin{aligned} I.\text{in}(n) + I_{12}.\text{out}(n) &= I.\text{in}(n) + I_1.\text{out}(n) + I_2.\text{out}(n) \\ &= I_{23}.\text{in}(n) + I_2.\text{out}(n) && (\text{As } I = I_{23} \oplus I_1) \\ &= I_3.\text{in}(n). && (\text{As } I_{23} = I_2 \oplus I_3) \end{aligned}$$

The condition on outflows is true because

$$\begin{aligned} I.\text{out}(n) &= I_1.\text{out}(n) + I_{23}.\text{out}(n) \\ &= I_1.\text{out}(n) + I_2.\text{out}(n) + I_3.\text{out}(n) \\ &= I_{12}.\text{out}(n) + I_3.\text{out}(n). \end{aligned}$$

Thus, if the LHS is defined and equal to I , then the RHS is defined and equal to I . By symmetry, and commutativity, the other direction is true as well.

- I_\emptyset is an identity with respect to \oplus :
This follows directly from the definitions.
- \oplus is cancellative, i.e. $I_1 \oplus I_2 = I_1 \oplus I_3 \Rightarrow I_2 = I_3$:
Let $I = I_1 \oplus I_2 = I_1 \oplus I_3$. Since the domains of I_2 and I_3 must be disjoint from I_1 and yet sum to the domain of I , they must be equal. Now for $n \in I_2$, by the definition of \oplus ,

$$I_2.in(n) = I.in(n) + I_1.out(n) = I_3.in(n),$$

so the inflows are equal. As for the outflows, if $n \notin I$ then

$$I.out(n) = I_1.out(n) + I_2.out(n) = I_1.out(n) + I_3.out(n),$$

which, by cancellativity of the flow domain, implies that $I_2.out(n) = I_3.out(n)$. On the other hand, if $n \in I_1$, then

$$I_1.in(n) = I.in(n) + I_2.out(n) = I.in(n) + I_3.out(n),$$

which again, by cancellativity, implies that $I_2.out(n) = I_3.out(n)$.

□

E.2 Section §4

Theorem 2 (Replacement Theorem). *If $I = I_1 \oplus I_2$, and $I_1 \lesssim I'_1$ are all valid interfaces such that $I'_1 \cap I_2 = \emptyset$ and $\forall n \in I'_1 \setminus I_1. I_2.out(n) = 0$, then there exists a valid $I' = I'_1 \oplus I_2$ such that $I \lesssim I'$.*

Proof. Let H, H_1, H_2 be flow graphs with interfaces I, I_1, I_2 respectively (thus, $H = H_1 \odot H_2$). Let,

$$H_0 = (N_0, e_0, fl_0) := (\text{dom}(I'_1) \setminus \text{dom}(I_1), (\lambda n, n'. \lambda_0), (\lambda n. in'_1(n) + out_1(n))).$$

First, we show that $H'_1 = H_0 \odot H_1$ is defined. The graph composition is defined because $I'_1 \cap I_2 = \emptyset$, and to show that it is a flow graph, we choose $in'_1 = I'_1.in$. We need to show, for any $n \in H'_1$, that

$$\begin{aligned} fl'_1(n) &= in'_1(n) + \sum_{n' \in H'_1} fl'_1(n') \triangleright e'_1(n', n) \\ &= in'_1(n) + \sum_{n' \in H_1} fl_1(n') \triangleright e_1(n', n). \quad (\text{as } n' \in H_0 \Rightarrow e'_1(n', n) = \lambda_0) \end{aligned}$$

When $n \in H_0$, by definition,

$$fl'_1(n) = fl_0(n) = in'_1(n) + out_1(n) = in'_1(n) + \sum_{n' \in H_1} fl_1(n') \triangleright e_1(n', n).$$

On the other hand, when $n \in H_1$,

$$fl'_1(n) = fl_1(n) = in_1(n) + \sum_{n' \in H_1} fl_1(n') \triangleright e_1(n', n) = in'_1(n) + \sum_{n' \in H_1} fl_1(n') \triangleright e_1(n', n)$$

($I_1 \lesssim I'_1$ implies $in_1(n) = in'_1(n)$).

Second, we show that $\text{int}(H'_1) = I'_1$. By construction, $\text{inf}(H'_1) = in'_1 = I'_1 \cdot \text{in}$. The outflow $\text{outf}(H'_1)(n)$ is equal to

$$\sum_{n' \in H'_1} fl'_1(n') \triangleright e'_1(n', n) = \sum_{n' \in H_1} fl_1(n') \triangleright e_1(n', n) = out_1(n) = out'_1(n)$$

(the last equality follows from $I_1 \lesssim I'_1$).

Third, we show that $H' = H'_1 \odot H_2 = H_0 \odot H$ is a flow graph. The graph composition is defined because $I'_1 \cap I_2 = \emptyset$, and to show that it is a flow graph, we choose the inflow

$$in'(n) = \begin{cases} in(n) & n \in H \\ in'_1(n) & n \in H_0. \end{cases}$$

We now need to show that (FlowEqn) holds, but first note that for any $n \in H'$,

$$\sum_{n' \in H'} fl'(n') \triangleright e'(n', n) = \sum_{n' \in H_1} fl_1(n') \triangleright e_1(n', n) + \sum_{n' \in H_2} fl_2(n') \triangleright e_2(n', n) \quad (1)$$

as when $n' \in H_0$, $e'(n', n) = e_0(n', n) = \lambda_0$. Now when $n \in H_0$,

$$fl'(n) = fl_0(n) = in'_1(n) + out_1(n) = in'(n) + \sum_{n' \in H_1} fl_1(n') \triangleright e_1(n', n).$$

The RHS of (FlowEqn) is, by (1),

$$in'(n) + \sum_{n' \in H_1} fl_1(n') \triangleright e_1(n', n) + \sum_{n' \in H_2} fl_2(n') \triangleright e_2(n', n)$$

but note that the third term is zero as $n \in I'_1 \setminus I_1 \Rightarrow I_2 \cdot \text{out}(n) = 0$. Thus, the flow equation holds. On the other hand, when $n \in H$,

$$fl'(n) = fl(n) = in(n) + \sum_{n' \in H} fl(n') \triangleright e(n', n) = in'(n) + \sum_{n' \in H'} fl'(n') \triangleright e'(n', n),$$

where the last equality follows from (1).

Finally, we let $I' = \text{int}(H')$ and show that $I \lesssim I'$. By construction, $\text{dom}(I) \subseteq \text{dom}(I')$ and $in(n) = in'(n)$ when $n \in \text{dom}(I)$. When $n' \notin \text{dom}(I)$, by (1), $out'(n') = out_1(n') + out_2(n') = out(n')$, completing the proof. \square

Before we prove Lemma 4, we prove some useful lemmas about the capacity and flow domains that are endomorphisms.

Lemma 7. *The capacity is equal to the following sum-of-paths expression:*

$$\text{cap}^i(G)(n, n') = \delta_{n=n'} + \sum_{\substack{n_1, \dots, n_k \in G \\ 0 \leq k < i}} e(n, n_1) \cdots e(n_k, n')$$

where when $k = 0$, $e(n, n_1) \cdots e(n_k, n')$ denotes $e(n, n')$.

Proof. Follows from the definition of capacity by a straightforward induction on i . \square

Lemma 8. *If $(M, +, 0, E)$ is a flow domain such that E is a closed set of endomorphisms, $G = (N, e)$ is a graph, $in: N \rightarrow M$ is an inflow such that $\text{FlowEqn}(in, e, fl)$, and $L \geq 1$,*

$$\begin{aligned} fl(n) &= in(n) + \sum_{\substack{n_1, \dots, n_k \in N \\ 1 \leq k < L}} in(n_1) \triangleright e(n_1, n_2) \cdots e(n_{k-1}, n_k) \triangleright e(n_k, n) \\ &+ \sum_{n_1, \dots, n_L \in N} fl(n_1) \triangleright e(n_1, n_2) \cdots e(n_{L-1}, n_L) \triangleright e(n_L, n). \end{aligned}$$

Proof. If $L = 1$, then this follows directly from (FlowEqn):

$$fl(n) = in(n) + \sum_{n' \in N} fl(n') \triangleright e(n', n).$$

For $L > 1$, by induction, we can assume that

$$\begin{aligned} fl(n) &= in(n) + \sum_{\substack{n_1, \dots, n_k \in N \\ 1 \leq k < L-1}} in(n_1) \triangleright e(n_1, n_2) \cdots e(n_{k-1}, n_k) \triangleright e(n_k, n) \\ &+ \sum_{n_1, \dots, n_{L-1} \in N} fl(n_1) \triangleright e(n_1, n_2) \cdots e(n_{L-2}, n_{L-1}) \triangleright e(n_{L-1}, n). \end{aligned}$$

Now, using (FlowEqn) to substitute for $fl(n_1)$, the third term becomes the following expression (after using the endomorphism property and some variable renaming):

$$\begin{aligned} &\sum_{n_1, \dots, n_{L-1} \in N} in(n_1) \triangleright e(n_1, n_2) \cdots e(n_{L-2}, n_{L-1}) \triangleright e(n_{L-1}, n) \\ &+ \sum_{n_1, \dots, n_L \in N} fl(n_1) \triangleright e(n_1, n_2) \cdots e(n_{L-1}, n_L) \triangleright e(n_L, n). \end{aligned}$$

Adding these to the above expression completes the proof. \square

Lemma 4. *If $(M, +, 0, E)$ is a flow domain such that M is a positive monoid and E is a nilpotent set of endomorphisms, then this flow domain has unique flows.*

Proof. Let $G = (N, e)$ be a graph, $in: N \rightarrow M$ an inflow, and $fl: N \rightarrow M$ an arbitrary flow that satisfies $\text{FlowEqn}(in, e, fl)$. By Lemma 8,

$$\begin{aligned} fl(n) &= in(n) + \sum_{\substack{n_1, \dots, n_k \in N \\ 1 \leq k < L}} in(n_1) \triangleright e(n_1, n_2) \cdots e(n_{k-1}, n_k) \triangleright e(n_k, n) \\ &+ \sum_{n_1, \dots, n_L \in N} fl(n_1) \triangleright e(n_1, n_2) \cdots e(n_{L-1}, n_L) \triangleright e(n_L, n). \end{aligned} \quad (2)$$

Let S be the third term in (2); we claim S is zero for sufficiently large L . When $L > p|N|$, where p is the degree of nilpotence of E , we will show that (in terms of the monoid order \leq):

$$S \leq \sum_{n_1, m \in N} fl(n_1) \triangleright \text{cap}^L(G)(n_1, m) \triangleright (\text{cap}^L(G)(m, m))^p \triangleright \text{cap}^L(G)(m, n).$$

If this is true, then since E is nilpotent, $(\text{cap}^L(G)(m, m))^p \equiv \lambda_0$ which implies that the right hand side is 0. But by positivity, $S \leq 0 \Rightarrow S = 0$. To show the above inequality, consider any term $fl(n_1) \triangleright e(n_1, n_2) \cdots e(n_{L-1}, n_L) \triangleright e(n_L, n)$ in S . As $L > p|N|$, by the Pigeonhole Principle there must be some $m \in N$ that appears $p + 1$ times in the sequence n_1, \dots, n_L . Thus, there exists some indices $1 \leq i_1 < \dots < i_{p+1} \leq L$ such that our term is equal to

$$\begin{aligned} & fl(n_1) \triangleright e(n_1, n_2) \cdots e(n_{i_1-1}, m) \\ & \triangleright \underbrace{(e(m, n_{i_1+1}) \cdots e(n_{i_2-1}, m)) \cdots (e(m, n_{i_p+1}) \cdots e(n_{i_{p+1}-1}, m))}_{p \text{ times}} \\ & \triangleright e(m, n_{i_{p+1}+1}) \cdots e(n_{L-1}, n_L) \triangleright e(n_L, n). \end{aligned}$$

By Lemma 7, this is $\leq fl(n_1) \triangleright \text{cap}^L(G)(n_1, m) \triangleright (\text{cap}^L(G)(m, m))^p \triangleright \text{cap}^L(G)(m, n)$. Thus,

$$fl(n) = in(n) + \sum_{\substack{n_1, \dots, n_k \in N \\ 1 \leq k < L}} in(n_1) \triangleright e(n_1, n_2) \cdots e(n_{k-1}, n_k) \triangleright e(n_k, n).$$

As this expression is uniquely determined by the flow domain, the inflow, and the graph, the flow must be unique. \square

To prove that EA graphs have unique flows, we first prove the following lemma: that the flow is equal to the capacity applied to the inflow.

Lemma 9. *If $(M, +, 0, E)$ is a flow domain such that M is a positive monoid, E is a closed set of endomorphisms, and $H = (N, e, fl)$ is an effectively acyclic flow graph, then*

$$fl(n) = \sum_{n' \in N} in(n') \triangleright \text{cap}(H)(n', n).$$

Proof. . \square

Corollary 1. *If $(M, +, 0, E)$ is a flow domain such that M is a positive monoid, E is a closed set of endomorphisms, and $H = (N, e, fl)$ is an effectively acyclic flow graph, then*

$$fl(n) = in(n) + \sum_{\substack{n_1, \dots, n_k \in N \\ 0 \leq k \leq |N|}} in(n_1) \triangleright e(n_1, n_2) \cdots e(n_k, n).$$

Lemma 5. *Let $(M, +, 0, E)$ be a flow domain such that M is a positive monoid and E is a closed set of endomorphisms. Given a graph (N, e) over this flow domain and inflow $in : N \rightarrow M$, if there exists a flow graph $H = (N, e, fl)$ that is effectively acyclic, then fl is unique.*

Proof. By Lemma 8, if we can show that for some $L > 1$ the following expression is 0, then the flow at every node is determined uniquely by the inflow and the graph:

$$\sum_{n_1, \dots, n_L \in N} fl(n_1) \triangleright e(n_1, n_2) \cdots e(n_{L-1}, n_L) \triangleright e(n_L, n). \quad (3)$$

Pick $L > |G|$. For any term T in (3), the Pigeonhole Principle tells us there must be some $m \in N$ that appears twice in the sequence n_1, \dots, n_L . Thus, there exists indices $1 \leq i_1 < i_2 \leq L$ such that

$$\begin{aligned} T = fl(n_1) \triangleright & \underbrace{e(n_1, n_2) \cdots e(n_{i_1-1}, m)}_{T_1} \triangleright \underbrace{e(m, n_{i_1+1}) \cdots e(n_{i_2-1}, m)}_{T_2} \\ & \triangleright \underbrace{e(m, n_{i_2+1}) \cdots e(n_{L-1}, n_L)}_{T_3} \triangleright e(n_L, n). \end{aligned}$$

Note that by Lemma 8, $fl(n_1) \triangleright T_1 \leq fl(m)$. Furthermore, as H is EA, $fl(m) \triangleright T_2 = 0$. Thus, $T = fl(n_1) \triangleright T_1 \triangleright T_2 \triangleright T_3 \leq fl(m) \triangleright T_2 \triangleright T_3 = 0 \triangleright T_3 = 0$. By positivity, $T = 0$, completing the proof. \square

Before we prove Theorem 3, we list some properties of flow graph compositions where one subgraph consists of disconnected nodes (all edges are zero function):

Lemma 10. *If $H' = H_0 \odot H$ such that $H_0 = (_, e_0, _)$ and $\forall n, n'. e_0(n, n') \equiv \lambda_0$, and H is EA, then H' is EA.*

Lemma 11. *If $H' = H_0 \odot H$ such that $H_0 = (_, e_0, _)$ and $\forall n, n'. e_0(n, n') \equiv \lambda_0$, then $\forall n' \notin H'. \text{outf}(H')(n') = \text{outf}(H)(n')$.*

Lemma 12. *If $H' = H_0 \odot H$ such that $H_0 = (_, e_0, _)$ and $\forall n, n'. e_0(n, n') \equiv \lambda_0$, then for all $n \in H', n' \notin H'$, and m ,*

$$m \leq \text{inf}(H')(n) \Rightarrow m \triangleright \text{cap}(H')(n, n') = \begin{cases} m \triangleright \text{cap}(H)(n, n') & n \in H \\ 0 & n \in H_0. \end{cases} \quad (4)$$

We now have enough tools to prove the analogue of the Replacement Theorem for effectively acyclic graphs.

Theorem 3. *Let $(M, +, 0, E)$ be a flow domain such that M is a positive monoid and E is a reduced set of endomorphisms. If $H = H_1 \odot H_2$ and $H_1 \lesssim_s H'_1$ are all effectively acyclic flow graphs such that $H'_1 \cap H_2 = \emptyset$ and $\forall n \in H'_1 \setminus H_1. \text{outf}(H_2)(n) = 0$, then there exists an effectively acyclic flow graph $H' = H'_1 \odot H_2$ such that $H \lesssim_s H'$.*

Proof. We first consider the case where $\text{dom}(H_1) = \text{dom}(H'_1)$. In this case, note that the definitions of contextual and subflow-preserving extensions reduce to $\text{int}(H_1) = \text{int}(H'_1)$ and $\forall n \in H_1, n' \notin H_1$,

$$m \leq \text{inf}(H_1)(n) \Rightarrow m \triangleright \text{cap}(H_1)(n, n') = m \triangleright \text{cap}(H'_1)(n, n'). \quad (5)$$

Thus, Lemma 2 tells us that $H' = H'_1 \odot H_2$ exists and $\text{int}(H') = \text{int}(H)$. All that remains to show is that H' is EA and show the analogue of (5) between H and H' . To simplify the proof, let $H'_2 := H_2$ and note that (5) also holds between H_2 and H'_2 .

Let us first show that H' is EA. If not, then there must exist some alternating sequence $n_1^1, n_1^2, \dots, n_1^{p_1}, n_2^1, \dots, n_k^{p_k} \in H$ such that n_i^1 and n_{i+1}^1 are in different subgraphs, and every n_i^j is in the same subgraph as n_i^1 , and some $1 \leq x \leq k, 1 \leq y \leq p_x$ such that

$$S = \text{in}(n_1^1) \triangleright e'(n_1^1, n_1^2) \cdots e'(n_k^{p_k-1}, n_k^{p_k}) \triangleright e'(n_k^{p_k}, n_x^y) \neq 0.$$

Note that k cannot be 1, because then the entire path would be in H'_i for $i \in \{1, 2\}$, contradicting the fact that H'_i is EA; thus $1 < k$.

If $x = k$, then let

$$T = \text{in}(n_1^1) \triangleright \left(\sum_{\substack{m_1, \dots, m_l \in H_- \\ 0 \leq l < L}} e'(n_1^1, m_1) \cdots e'(m_l, n_2^1) \right) \cdots \\ \left(\sum_{\substack{m_1, \dots, m_l \in H_- \\ 0 \leq l < L}} e'(n_{k-1}^1, m_1) \cdots e'(m_l, n_k^1) \right)$$

for some L such that $S + _ = T \triangleright e'(n_k^1, n_k^2) \cdots e'(n_k^{p_k}, n_x^y)$. As $S \neq 0$, positivity tells us $S + _ \neq 0$, and so $T \neq 0$. By applying (5) repeatedly, we can show that $T = \text{in}(n_1^1) \triangleright c_(n_1^1, n_2^1) \cdots c_(n_{k-1}^1, n_k^1)$. The definition of interface composition tells us that $\text{in}_(n_k^1) = T + _$, so by positivity we must have $\text{in}_(n_k^1) \triangleright e'(n_k^1, n_k^2) \cdots e'(n_k^{p_k}, n_x^y) \neq 0$. However, this is a path entirely in some subgraph, contradicting EA of that subgraph.

On the other hand, if $x < k$, we have

$$S = \text{in}(n_1^1) \triangleright e'(n_1^1, n_1^2) \cdots e'(n_x^{y-1}, n_x^y) \triangleright (e'(n_x^y, n_x^{y+1}) \cdots e'(n_x^{p_x}, n_{x+1}^1)) \\ \triangleright e'(n_{x+1}^1, n_{x+1}^2) \cdots e'(n_k^{p_k}, n_x^y).$$

Since our flow domain is reduced and $S \neq 0$, we can repeat the parenthesized portion and obtain

$$S' = \text{in}(n_1^1) \triangleright e'(n_1^1, n_1^2) \cdots e'(n_x^{y-1}, n_x^y) \triangleright e'(n_x^y, n_x^{y+1}) \cdots e'(n_x^{p_x}, n_{x+1}^1) \\ \triangleright e'(n_{x+1}^1, n_{x+1}^2) \cdots e'(n_k^{p_k}, n_x^y) \triangleright e'(n_x^y, n_x^{y+1}) \cdots e'(n_x^{p_x}, n_{x+1}^1) \neq 0.$$

Note that this is a path that starts and ends on the boundary (i.e. with the last edge crossing between subgraphs). Depending on whether n_x^y and n_k^1 are in the same subgraph or not, we can use new variable names and write this as

$$S' = \text{in}(m_1^1) \triangleright e'(m_1^1, m_1^2) \cdots e'(m_l^{q_l-1}, m_l^{q_l}) \triangleright e'(m_l^{q_l}, m_{l+1}^1) \neq 0.$$

Let

$$T = in(m_1^1) \triangleright \left(\sum_{\substack{n_1, \dots, n_k \in H_- \\ 0 \leq k < L}} e'(m_1^1, n_1) \cdots e'(n_k, m_2^1) \right) \cdots \\ \left(\sum_{\substack{n_1, \dots, n_k \in H_- \\ 0 \leq k < L}} e'(m_l^1, n_1) \cdots e'(m_k, n_{l+1}^1) \right)$$

for some L such that $T = S' + _$, so by positivity, $T \neq 0$. Again, applying (5) repeatedly lets us show that $T = in(m_1^1) \triangleright c_-(m_1^1, m_2^1) \cdots c_-(m_l^1, m_{l+1}^1)$. By expanding the definitions of the capacities and using the endomorphism property we get that T is the sum of cyclic paths in H , at least one of which is non-zero. This contradicts the fact that H is EA. The proof of the analogue of (5) between H and H' proceeds similarly, by breaking up all paths into segments in each subgraph and using (5) repeatedly; we omit it here for clarity of presentation.

We now turn to the case where $\text{dom}(H_1) \subset \text{dom}(H'_1)$. Our aim is to reduce this case to the previous case, thus we need to construct an intermediate graph H''_1 that satisfies the conditions of this theorem but additionally has the same domain as H'_1 .

We define $H_0 = (N_0, e_0, fl_0) := (\text{dom}(H'_1) \setminus \text{dom}(H_1), (\lambda n, n'. \lambda_0), (\lambda n. in'_1(n) + out_1(n)))$. Let $H''_1 = H_0 \odot H_1$; this exists because the graphs are disjoint by construction, and we show that (FlowEqn) is true for inflow in'_1 by showing, for any $n \in H''_1$, that

$$fl''_1(n) = in'_1(n) + \sum_{n' \in H''_1} fl''_1(n') \triangleright e''_1(n', n) \\ = in'_1(n) + \sum_{n' \in H_1} fl_1(n') \triangleright e_1(n', n). \quad (\text{as } n' \in H_0 \Rightarrow e''_1(n', n) = \lambda_0)$$

When $n \in H_0$, by definition,

$$fl''_1(n) = fl_0(n) = in'_1(n) + out_1(n) = in'_1(n) + \sum_{n' \in H_1} fl_1(n') \triangleright e_1(n', n).$$

On the other hand, when $n \in H_1$,

$$fl''_1(n) = fl_1(n) = in_1(n) + \sum_{n' \in H_1} fl_1(n') \triangleright e_1(n', n) = in'_1(n) + \sum_{n' \in H_1} fl_1(n') \triangleright e_1(n', n),$$

where we use $in_1(n) = in'_1(n)$ from $H_1 \lesssim_s H'_1$. By Lemma 10, we get that H''_1 is EA.

Next, we show that $H''_1 \lesssim_s H'_1$. First, to show that $I''_1 \lesssim I'_1$, note that by construction $\text{dom}(H''_1) = \text{dom}(H'_1)$; we have already shown that H''_1 satisfies the flow equation with inflow in'_1 , hence they have the same inflows; and by Lemma 11 they have the same outflows. Second, let $n \in H''_1, n' \notin H'_1$ and $m \leq \text{inf}(H''_1)(n)$. If $n \in H_0$, then by Lemma 12, $m \triangleright \text{cap}(H''_1)(n, n') = 0$ and since $n \notin H_1$, by $H_1 \lesssim_s H'_1$, $m \triangleright \text{cap}(H'_1)(n, n') = 0$. On the other hand, if $n \in H_1$, then by Lemma 12, $m \triangleright$

$\text{cap}(H_1'')(n, n') = m \triangleright \text{cap}(H_1)(n, n')$ which equals $m \triangleright \text{cap}(H_1')(n, n')$ by $H_1 \lesssim_s H_1'$ and $\text{inf}(H_1'')(n) = \text{inf}(H_1')(n) = \text{inf}(H_1)(n)$. The third condition of $H_1'' \lesssim_s H_1'$ is vacuously true since $\text{dom}(H_1'') = \text{dom}(H_1')$.

By a similar argument, we can show that $H'' = H_1'' \odot H_2 = H_0 \odot H$ exists, and is EA. At this point, we have $H'' = H_1'' \odot H_2$ and $H_1'' \lesssim_s H_1'$ are all effectively acyclic flow graphs such that $H_1' \cap H_2 = \emptyset$ and additionally $\text{dom}(H_1'') = \text{dom}(H_1')$. Thus, we can use the case when the domains are equal (which has been proved above) and obtain an EA flow graph $H' = H_1' \odot H_2$ such that $H'' \lesssim_s H'$. Since \lesssim_s is transitive, we just need to show that $H \lesssim_s H''$ to complete the proof.

To see that $H \lesssim_s H''$, first consider $I \lesssim I''$: by construction and $H_1 \lesssim_s H_1'$, $\text{dom}(H) \subseteq \text{dom}(H'')$; we constructed H'' to have an inflow that is equal to in on all nodes in H ; and by Lemma 11 they have the same outflows. Second, let $n \in H, n' \notin H''$ and $m \leq \text{inf}(H)(n)$. By Lemma 12, $m \triangleright \text{cap}(H'')(n, n') = m \triangleright \text{cap}(H)(n, n')$. Third, let $n \in H'' \setminus H, n' \notin H''$ and $m \leq \text{inf}(H'')(n)$. Since $n \in H'' \setminus H \Rightarrow n \in H_0$, Lemma 12 tells us that $m \triangleright \text{cap}(H'')(n, n') = 0$. \square