The Ramifications of Sharing in Data Structures

Aquinas Hobor
National University of Singapore
hobor@comp.nus.edu.sg

Jules Villard
University College London
j.villard@cs.ucl.ac.uk

Abstract

Programs manipulating mutable data structures with intrinsic sharing present a challenge for modular verification. Deep aliasing inside data structures dramatically complicates reasoning in isolation over parts of these objects because changes to one part of the structure (say, the left child of a dag node) can affect other parts (the right child or some of its descendants) that may point into it. The result is that finding intuitive and compositional proofs of correctness is usually a struggle. We propose a compositional proof system that enables local reasoning in the presence of sharing.

While the AI "frame problem" elegantly captures the reasoning required to verify programs without sharing, we contend that natural reasoning about programs with sharing instead requires an answer to a different and more challenging AI problem, the "ramification problem": reasoning about the indirect consequences of actions. Accordingly, we present a RAMIFY proof rule that attacks the ramification problem head-on and show how to reason with it. Our framework is valid in any separation logic and permits sound compositional and local reasoning in the context of both specified and unspecified sharing. We verify the correctness of a number of examples, including programs that manipulate dags, graphs, and overlaid data structures in nontrivial ways.

Categories and Subject Descriptors F.3.1 [*Specifying and Verifying and Reasoning about Programs*]: Logics of programs; D.2.4 [*Software/Program Verification*]: Correctness proofs, Formal methods

General Terms Languages, Theory, Verification.

Keywords Aliasing, Heap/Shape, Modularity, Separation logic.

1. Introduction

Data structures with intrinsic sharing, such as acyclic and unrestricted graphs as well as various kinds of overlaid data structures, are pervasive in computing. An example of an overlaid data structure can be found in the Linux deadline I/O scheduler, in which the set of events forms both a singly linked list and a binary sorted tree, depending on which links one follows. Programs manipulating data structures with sharing are often short, but the reason that they are correct can be subtle, and previous work has not come up with general, intuitive and compositional principles for reasoning about such programs. The key difficulty is that deep aliasing dramatically complicates reasoning in isolation over parts of these objects: changes to one part of the structure (say, the left child of a dag) can affect other parts (the right child or its descendants) that may point into it.

We propose a compositional proof system for programs manipulating shared data structures. Our framework directly addresses the intrinsic sharing present in the data structures and achieves compositionality via applications of the following *ramify rule*:

$$\textsc{Ramify}\quad \frac{\{P\}\ c\ \{Q\} \qquad ramify(R,P,Q,R')}{\{R\}\ c\ \{R'\}}$$

At first glance there seems to be no connection between the known spec $\{P\}\ c\ \{Q\}$ and the desired spec $\{R\}\ c\ \{R'\}$. The connection is given by the *ramification*, indicated by the $ramify(R,P,Q,R')$ premise, which asserts (semantically, although this paper also provides ways to reason syntactically about it) that the "global" assertion R becomes R' after a "local" transformation from P to Q.

The term "ramification" comes from artificial intelligence [Fin87, Thi01] and refers to the problem of understanding the indirect (global) consequences of (local) actions (*e.g.* relocating a bookcase might reduce the ambient light by blocking the window). Ramification is contrasted with the simpler "frame" problem, which centers on maintaining knowledge after unrelated actions (*e.g.*, relocating the bookcase does not change the number of moons of Jupiter).

Program verification has had significant success handling the frame problem, especially with the frame rule of separation logic [Rey02]:

$$\textsc{Frame}\quad \frac{\{P\}\ c\ \{Q\}}{\{P * F\}\ c\ \{Q * F\}}$$

Here the *separating conjunction* $*$ ensures that P and F cover *disjoint* pieces of heap, allowing the frame rule to guarantee that F is unchanged under the action of c. The frame rule buys us compositionality in the presence of the heap: we can reason about the effect a program has on the portions of heap it accesses, and reuse that spec in any bigger heap. This has given rise to concise, compositional proofs of programs, even in the presence of *some* forms of sharing where one knows *what* is shared *by whom*.

Unfortunately, we usually cannot use the frame rule directly when verifying programs that manipulate data structures with *unrestricted* sharing because such structures cannot easily be massaged into the form $P * F$: for example, the left and right descendants of a dag node are not usually disjoint. The reason to focus on ramification rather than frame is that the former allows us to reuse specs for c in far more diverse settings than the latter permits. Of course, with great power comes great responsibility: having isolated the parts of the proof that require careful examination of indirect effects on the global structure, we are left with ramification obligations to prove.

As it turns out, ramifications are expressible as separation logic entailments: $ramify(R,P,Q,R') \stackrel{\text{def}}{=} R \vdash P * (Q \mathrel{-\!\!*} R')$. These entailments feature the "magic wand" connective of separation logic ("for all states σ_1 satisfying Q and disjoint from the current state σ_2, the combination of both states $\sigma_1 \oplus \sigma_2$ satisfies R'"), which is notor-

POPL'13 January 23–25, 2013, Rome, Italy.

iously hard to reason about in general given the universal quantification over states. However, appearances of $-\!\!*$ in ramifications are restricted to a particular idiom that, together with $*$, denotes an *update* to the state. Guided by this intuition, we are able to reduce these spatial entailments to more abstract reasoning about the nature of the update on the structure's mathematical representation (*e.g.*, graphs as sets of nodes and edges and transformations on said graphs). The verification process thus divides into two parts: first, showing that a concrete program correctly implements some transformation on an abstract mathematical structure; and second, showing that those mathematical transformations produce the desired specification.

This division gives us the freedom to describe data structures with intrinsic sharing in the most natural way. We will present examples that use the separating conjunction $*$ of separation logic to reason about genuine disjointness (*e.g.*, between the parent of a dag node and its children), the overlapping conjunction $\circledast$ to reason about unspecified sharing (*e.g.*, between the left and right children of a dag node), and the classical conjunction to reason about complete sharing (*e.g.*, an overlaid data structure).

In contrast to previous work, we achieve *compositional* reasoning and *embrace* the sharing. Approaches based on separation logic favored convoluted invariants that hacked the state into the disjoint pieces required by the frame rule. Often the predicate definitions depended heavily on the program at hand (*e.g.* the dag definition used could depend on the order of traversal in the algorithm [BCO04]). In other words, previous attempts to reason about shared data structures with separation logic have stood on their head to avoid the sharing. Other approaches suffered from these problems at least as much and often gave up compositionality altogether [Bor00].

Our key contributions are as follows:

- We present the RAMIFY rule which enables local reasoning while accounting for global effects precisely when they are required. Ramification can reason about programs that manipulate data structures with unrestricted sharing while enabling the small specifications, compositionality, and expressiveness that have led to separation logic's success.
- Although the ramify rule leads to more natural Hoare proofs, the entailment checks can be nontrivial. We have developed a "ramification library" of lemmas that help simplify the ramification conditions. Crucially, we also show how to *prove* ramifications concerned with certain general graph and dag updates in a way that enables a separation of concern between heap manipulations and mathematical reasoning about graphs.
- We have applied the ramify rule to a variety of algorithms that manipulate data structures with nontrivial sharing. Although some of the examples are not long, all involve intricate reasoning due to the heavy use of sharing. We think that a strength of our approach is that the Hoare invariants at each program point are natural and seem to follow our "programmer's intuition" much more closely than traditional proofs.
- We give a semantic account of ramification, and show that RAMIFY and FRAME are each derivable from the other, meaning that our framework is applicable in any separation logic. Moreover, we identify the precise constraints on the underlying model that enable the overlapping conjunction $\circledast$, and show that most separation logics in the literature can therefore follow our recipe and use it to reason about unspecified sharing.

The rest of the paper is organized as follows: we first recall some important concepts from separation logic (§2). We then motivate and present the ramify rule (§3), and show how to reason about it (§4). Based on this, we provide proof sketches for three examples that showcase different aspects of ramification: marking a dag (§5), removing from an overlaid data structure (§6), and Cheney's garbage collector (§7). Finally, we show how ramification is applicable in virtually any separation logic (§8), compare to related works, and conclude.

2. Separation Logic and Trees

Recall the framework of separation logic [IO01, Rey02] while considering the following `mark` procedure, written in C, that recursively marks binary trees, dags, or graphs:

```
1 struct node {int m; struct node *l,*r;};
2 void mark(struct node *x) {
3   if (!x || x->m) return;
4   struct node *l = x->l, *r = x->r;
5   x->m = 1; mark(l); mark(r); }
```

Separation logic allows straightforward inductive definitions of predicates to describe tree-like data structures in the heap. The following definition disregards the actual contents and location of each node, but does make sure that the structure is acyclic (thanks to the $*$ between the root and the subtrees) and that no sharing occurs between subtrees (thanks to the $*$ between the children):

$$\mathsf{tree}(x) \stackrel{\text{def}}{=} (x = 0 \wedge \mathsf{emp}) \vee \exists d, l, r.\, x \mapsto d, l, r * \mathsf{tree}(l) * \mathsf{tree}(r)$$

The definition of `tree` uses the standard *classical* separation logic operators. A heaplet h satisfies the points-to predicate $x \mapsto y$ when h contains *only* the location x, whose value is y, and the separating conjunction $P * Q$ asserts that P and Q hold on disjoint subheaps. We use $x \mapsto d, l, r$ as a shorthand for $(x + 0) \mapsto d * (x + 1) \mapsto l * (x + 2) \mapsto r$ (simplifying the memory model so that *e.g.*, each datum occupies one unit of space).

It is well-known how to use separation logic to prove the `mark` procedure memory safe for trees. Moreover, the separation logic proof mirrors the programmer's intuitions beautifully. The crux of the verification is to handle the recursive calls via the frame rule, *e.g.*, at line 5, taking the spec of `mark` as a premise:

$$\frac{\{\mathtt{tree(l)}\}\ \mathtt{mark(l)}\ \{\mathtt{tree(l)}\}}{\{\mathtt{t} \mapsto 1, \mathtt{l}, \mathtt{r} * \mathtt{tree(l)} * \mathtt{tree(r)}\}\ \mathtt{mark(l)}\ \{\mathtt{t} \mapsto 1, \mathtt{l}, \mathtt{r} * \mathtt{tree(l)} * \mathtt{tree(r)}\}}\ \textsc{Frame} \tag{1}$$

This is a canonical example of how inductive predicates, the separating conjunction, and the frame rule fit together to produce concise proofs. Unrolling the `tree` predicate yields $*$-conjoined formulas, so the proof system, via its frame rule, is able to perform surgery on the symbolic state and work on each sub-state independently.

3. Ramifications for Sharing

We now turn to the case of data structures with sharing, and introduce our RAMIFY rule. We begin by defining inductive predicates for dags and graphs before presenting the proof sketch that we aspire to for the `mark` procedure when applied to dags.

3.1 Dag and Graph Predicates

Our first task is to define a dag predicate. Since the separating conjunction $*$ prevents sharing, our first attempt updates `tree` to utilize regular conjunction $\wedge$ between the children instead:

$$\mathsf{dag}_0(x) \stackrel{\text{def}}{=} (x = 0 \wedge \mathsf{emp}) \vee \exists l, r.\, x \mapsto l, r * (\mathsf{dag}_0(l) \wedge \mathsf{dag}_0(r))$$

Unfortunately, in *classical* separation logic, $\mathsf{dag}_0(x)$ actually describes a linked list because the conjunction forces the two sub-dags to occupy *exactly* the same space in memory ($h \models P \wedge Q$ if $h \models P$ and $h \models Q$). However, Reynolds points out that dag_0 is correct in *intuitionistic* separation logic, in which $x \mapsto y$ holds on any heap that contains *at least* x, rather than *only* x [Rey02, §6]. Translated into our classical setting this is equivalent to defining dags as follows:

$$\mathsf{dag}_1(x) \stackrel{\text{def}}{=} (x = 0 \wedge \mathsf{emp}) \vee \exists l, r.\, x \mapsto l, r * ((\mathsf{dag}_1(l) * \mathsf{true}) \wedge (\mathsf{dag}_1(r) * \mathsf{true}))$$

If our first attempt was in some sense "too small", then our second is "too big": $\mathsf{dag}_1(x)$ holds on any heap that contains *at least* a dag rooted at x. As usual in intuitionistic separation logic, it is impossible to fully verify certain algorithms (*e.g.* showing that dag disposal completely frees the structure) using dag_1.

What we want is a way to get the overlapping features of the intuitionistic conjunction without actually becoming intuitionistic. We turn to another connective, scarcely studied in the published literature, which we dub the *overlapping conjunction* and write $\circledast$, and which precisely characterizes the desired sharing:

$$h \models P \circledast Q \stackrel{\text{def}}{=} \exists h_1, h_2, h_3.\\ (h_1 \oplus h_2 \oplus h_3 = h) \wedge (h_1 \oplus h_2 \models P) \wedge (h_2 \oplus h_3 \models Q)$$

The $\oplus$ is the combination operator on the underlying separation algebra [COY07] (often some kind of disjoint union). Contrast the definition of $\circledast$ with the standard definition of $*$:

$$h \models P * Q \stackrel{\text{def}}{=} \exists h_1, h_2.\, (h_1 \oplus h_2 = h) \wedge (h_1 \models P) \wedge (h_2 \models Q)$$

Here are some properties of $\circledast$ for reference and to aid intuition.

Lemma 3.1

$$P \circledast \mathsf{emp} \quad \dashv\vdash \quad P \tag{2}$$
$$P \wedge Q \quad \vdash \quad P \circledast Q \tag{3}$$
$$P * Q \quad \vdash \quad P \circledast Q \tag{4}$$
$$P \circledast Q \quad \vdash \quad P * \mathsf{true} \tag{5}$$
$$P \circledast Q \quad \dashv\vdash \quad \exists R.\, (R \mathrel{-\!\!*} P) * (R \mathrel{-\!\!*} Q) * R \tag{6}$$
$$P \circledast Q \quad \dashv\vdash \quad Q \circledast P \tag{7}$$
$$P \circledast (Q \circledast R) \quad \dashv\vdash \quad (P \circledast Q) \circledast R \tag{8}$$
$$\mathit{covar}(F_1) \Rightarrow \mathit{covar}(F_2) \Rightarrow \mathit{covar}(\lambda P.\, F_1(P) \circledast F_2(P)) \tag{9}$$

Equations (2), (3), (4) and (5) are immediate from the definition of $\circledast$. We use quantification over predicates in (6). Commutativity (7) is direct from (6) and the commutativity of $*$. In contrast, associativity (8) is trickier and requires cross split (see §8.3). Finally, Lem. (9) enables $\circledast$ to be used in covariant[1] recursive predicates, just like $*$ and $\wedge$. Whenever we write recursive definitions using $\circledast$, including dag and graph below, we are implicitly using (9).

The key point to $\circledast$ is that we can use it in exactly the same places that feature the kinds of sharing that the intuitionistic $\wedge$ captures, but it does not "over-approximate" the resulting structure. That is, it allows us to define a classical dag (with a data field) as:

$$\mathsf{dag}(x) \stackrel{\text{def}}{=} (x = 0 \wedge \mathsf{emp}) \vee \\ \exists d, l, r.\, x \mapsto d, l, r * (\mathsf{dag}(l) \circledast \mathsf{dag}(r)) \tag{10}$$

The separating conjunction $*$ between the root x and its children prevents cycles in the data structure. Pleasingly, the definition for graphs simply replaces this remaining $*$ with another $\circledast$:

$$\mathsf{graph}(x) \stackrel{\text{def}}{=} (x = 0 \wedge \mathsf{emp}) \vee \\ \exists d, l, r.\, x \mapsto d, l, r \circledast \mathsf{graph}(l) \circledast \mathsf{graph}(r) \tag{11}$$

We will equip dag and graph with mathematical dags δ and graphs γ to enable proofs of functional correctness, writing $\mathsf{dag}(x, \delta)$ and $\mathsf{graph}(x, \gamma)$ respectively; δ and γ need not be "tight" and so can include vertices that are unreachable from x. Mathematical trees lack sharing and are hence directly definable as terms; mathematical dags δ and graphs γ are more complicated and so we defer the associated formal definitions until §4.2; one key notation is $\delta(x) = (d, l, r)$, which indicates that the mathematical node x is associated with data d and successors l and r.

Unspecified Sharing Observe that $\circledast$ models *unspecified sharing*: *i.e.*, the dag predicate does not say which parts of a dag are shared. In contrast, *specified sharing* requires the precise identification of the shared part, *e.g.* on a dag identifying which nodes are shared between the left and right children; often this is very difficult.

[1] $\mathit{covar}(F) \stackrel{\text{def}}{=} (P \vdash Q) \Rightarrow F(P) \vdash F(Q)$

```
1  void mark(struct node *x) { // {dag(x, δ)}
2    struct node *l,*r;
3    if (x == 0 || x->m == 1) return;
4    l = x->l; r = x->r;
5  // {x ↦ 0,l,r * (dag(l, δ) ⊛ dag(r, δ)) ∧ δ(x) = (0,l,r)}
6    x->m = 1;
7  // {x ↦ 1,l,r * (dag(l, δ) ⊛ dag(r, δ)) ∧ δ(x) = (0,l,r)}
8    mark(l);
9  //↯(23) { x ↦ 1,l,r * (dag(l, m(δ,l)) ⊛ dag(r, m(δ,l))) ∧
           δ(x) = (0,l,r) }
10   mark(r);
11 //↯(24) { x ↦ 1,l,r * (dag(l, δ') ⊛ dag(r, δ')) ∧
           δ(x) = (0,l,r) ∧ δ' = m(m(δ,l),r) }
12 } // {dag(x, m(δ, x))}
```

Figure 1: Proof sketch for marking a binary dag. The steps that induce ramifications are indicated with $\lightning_i$, where the associated ramification entailment is equation number i.

On the other hand, sometimes specified sharing is exactly what the doctor ordered. Although the overlapping conjunction is extremely useful, our framework is not based around it, and one of our key contributions is that RAMIFY can handle both specified and unspecified sharing. For an example of specified sharing, see §6, which uses $\wedge$ instead of $\circledast$; moreover, see §8.3 for how we can use the explicit overlapping conjunction of Cherini and Blanco [CB09].

3.2 Ramifications of Manipulating Dags

Fig. 1 presents the annotated[2] proof sketch of the functional correctness of `mark` when applied to dags using the small spec $\{\mathsf{dag}(\mathtt{l}, \delta)\}$ `mark(l)` $\{\mathsf{dag}(\mathtt{l}, m(\delta, \mathtt{l}))\}$. The function $m(\delta, x)$, whose formal definition is deferred until §5, indicates the mathematical dag derived from δ via marking starting from node x. Notice that this specification immediately implies that if the initial dag is unmarked then the final dag is completely marked.

As is the case for many recursive programs on graph-like data structures, part of the state tracking the recursive exploration of the graph resides in the call stack, which remembers which states have been only partially processed. Our spec accounts for this complexity while remaining local (*i.e.*, it only describes the portion of memory accessed by `mark`), enabling compositional reasoning. Moreover, we enjoy straightforward invariants at each program point.

Although the invariants are natural, the proof in separation logic is far from obvious. Things are straightforward enough until we reach the first recursive call at line 8. For tree we applied the frame rule in equation 1, which worked very well. While we can easily frame away the $\mathtt{x} \mapsto 1, \mathtt{l}, \mathtt{r}$ from the precondition (line 7), disentangling the two dag predicates into $\mathsf{dag}(\mathtt{l}, \delta)$ on the one hand and a $*$-disjoint frame on the other would necessitate describing the shape of the right child once everything that is shared with the left child has been removed, which is exactly what we wish to avoid. The second recursive call, in line 10, presents exactly the same problem: we wish to frame but cannot. These two recursive calls require a new proof pattern that we call *ramification*.

3.3 The RAMIFY Rule

While the proof outline of Fig. 1 provides all the invariants needed to prove `mark` on dags, FRAME cannot be applied directly to reason about the effect of applying the `mark` spec on the left child because the left and right child are not disjoint. To solve this issue, we introduce the *ramification rule*, which allows the reasoning to

[2] We often write *e.g.* $\delta(\mathtt{x}) = \ldots$ when what we really mean is $\mathtt{x} \Downarrow v \wedge \delta(v) = \ldots$, where $\mathtt{x} \Downarrow v$ means that the variable x evaluates to the value v in the current state, because mathematical graphs take values rather than variables. We elide these kinds of details for the presentation.

progress through commands that have indirect global effects:

$$\textsc{Ramify}\quad \frac{\{P\}\ c\ \{Q\} \qquad R \vdash P * (Q \mathrel{-\!\!*} R')}{\{R\}\ c\ \{R'\}}\ \begin{array}{l} fv(Q \mathrel{-\!\!*} R') \cap \\ modif(c) = \varnothing \end{array}$$

RAMIFY isolates the complicated leap in reasoning at each recursive call site so that the assertions at each program point remain natural, such as in Fig. 1 (*e.g.*, the assertions are free from $\mathrel{-\!\!*}$). No free variables of $Q \mathrel{-\!\!*} R'$ may be modified by c. As usual, magic wand (separating implication) is the adjunct[3] of $*$:

$$h \models P \mathrel{-\!\!*} Q \ \stackrel{\text{def}}{=}\ \forall h'.\, h \perp h' \ \Rightarrow\ h' \models P \Rightarrow h \oplus h' \models Q$$

Here $h \perp h'$ asserts the *compatibility* of h and h' ($\exists h''.\, h \oplus h' = h''$). Informally, ramify can be read as "the result of applying c in a state R is R' if replacing P inside R with Q yields R'". Magic wand binds *more loosely* than any other operator.

Ramify is sufficiently abstract that it can be hard to appreciate. As an initial demonstration of its power, observe that the frame rule (modulo some restrictions on free variables as discussed below) is a direct consequence because $P * F \vdash P * (Q \mathrel{-\!\!*} Q{*}F)$.

Next, let us apply ramification to verify the following spec, in which $\mathtt{x} \mapsto -$ is the standard notation for $\exists x'.\, \mathtt{x} \mapsto x'$.

$$\{\mathtt{x} \mapsto - \circledast \mathtt{y} \mapsto -\}\ \mathtt{*x\ =\ a}\ \{\mathtt{x} \mapsto \mathtt{a} \circledast \mathtt{y} \mapsto -\}$$

RAMIFY emits two subgoals. The first precisely matches the standard small axiom for store update in separation logic:

$$\{\mathtt{x} \mapsto -\}\ \mathtt{*x\ =\ a}\ \{\mathtt{x} \mapsto \mathtt{a}\}$$

The second is the following ramification entailment, whose proof is direct from the associated definitions:

$$x \mapsto - \circledast y \mapsto - \vdash x \mapsto - * (x \mapsto a \mathrel{-\!\!*} x \mapsto a \circledast y \mapsto -) \quad (12)$$

Free variables Notice that RAMIFY has a side condition to pacify the usual free variable bugaboo. Usually this is no big deal, but it causes trouble when we want to use ramification to verify commands of the form $\mathtt{x\ =} f(\ldots)$, since x *is* modified and we may want to refer to it in the postcondition. One sufficient solution, which pleasingly removes all free variable side conditions, is to use variables as resource [BCY06], but this introduces other complications. Another solution is to use the following variant of the ramify rule:

$$\textsc{RamifyAssign}\quad \frac{\{P\}\ x' = f(\ldots)\ \{Q\} \qquad R \vdash P * (Q \mathrel{-\!\!*} [x \mapsto x']R')}{\{R\}\ x = f(\ldots)\ \{R'\}}\ {}^{\dagger}$$

$${}^{\dagger} x' \notin fv(R, R', P) \cup vars(f)\ \wedge\ fv(Q \mathrel{-\!\!*} R') \cap modif(f) = \varnothing$$

RAMIFYASSIGN is a consequence of RAMIFY and the usual rules for assignment and sequence if we are allowed to make the local program transformation from $\mathtt{x\ =} f(\ldots)$ to $x'\mathtt{=}f(\ldots)\mathtt{;\ x\ =}x'$ in which x' is always chosen fresh. From now on we will sweep free variable issues under the rug, silently using RAMIFYASSIGN when needed.

Lookup Because points-to facts may be buried inside shared parts of the state, we find it convenient to use the *global* rule for lookup [Rey02] instead of the standard *local* one of separation logic:

$$\textsc{Lookup}\quad \frac{}{\{\exists x'.\, (y \mapsto x' * \mathsf{true}) \wedge [x \mapsto x']P\}\ \mathtt{x\ =\ *y}\ \{P\}}\ x' \notin fv(P, y)$$

In fact, RAMIFYASSIGN is able to derive LOOKUP from the standard local separation logic axiom.

4. Reasoning about Ramifications

To set the stage for the verification of our examples, we now present techniques for general reasoning about ramifications and link abstract mathematical reasoning about graphs to spatial ramifications.

[3] That is, $*$ and $\mathrel{-\!\!*}$ are related by $P * Q \vdash R \ \Leftrightarrow\ P \vdash Q \mathrel{-\!\!*} R$.

4.1 Ramification Library

Our *ramification library* is a collection of lemmas that help reduce complicated ramifications and related entailments. Some of the more general-purpose lemmas, which can handle simplifications such as removing frames that occur within ramifications, are grouped in Fig. 2. Other lemmas in our library are specific to certain data structures such as graphs; we will meet some of these in §4.3.

Some of the lemmas in Fig. 2 require that various predicates be *precise*, which means that whenever P is satisfied on a sub-state ($h_1 \leqslant h_3 \stackrel{\text{def}}{=} \exists h_2.\, h_1 \oplus h_2 = h_3$), that sub-state must be unique:

$$\begin{aligned} precise(P) \stackrel{\text{def}}{=} \forall h_1, h_2, h_3.\, h_1 \leqslant h_3 \Rightarrow h_2 \leqslant h_3 \Rightarrow \\ h_1 \models P \Rightarrow h_2 \models P \Rightarrow h_1 = h_2 \end{aligned}$$

All the predicates we consider here are indeed precise, so this is never a concern in this paper.

Lemmas 4.3 and 4.4 use $\mathrel{-\!\!\circledast}$, the existential magic wand:

$$h \models P \mathrel{-\!\!\circledast} Q \ \stackrel{\text{def}}{=}\ \exists h'.\, h \perp h' \wedge h' \models P \wedge h \oplus h' \models Q$$

This operator can be tricky because one does not know *which* copy of P has been pulled out of Q, but is handy sometimes.

Lem. 4.1 allows one to frame away F *within* an overlapping conjunction in order to focus on an easier entailment. Perhaps surprisingly, both P and Q need to be precise for it to hold.

Lem. 4.2 and 4.3 are analogues of Lem. 4.1 for classical conjunctions. Although Lem. 4.2 is more immediate, the premise $R \vdash P * (Q \mathrel{-\!\!*} R')$ may sometimes be difficult to establish. In particular, one would need to show that P can always be found in a state satisfying R, which is not necessarily the case; it is true in the original ramification because the state satisfies $(P * F) \wedge R$. Lem. 4.3 remedies this by requiring the subtler condition that *if* P may be found in R, then *any* way of adding Q instead of P yields R'. We use this lemma to prove the `pop` program in §6.

Lem. 4.4 allows one to ignore parts of the state that remain invariant during an update. For instance, a procedure may require some piece of state in its precondition that is not modified and thus percolates unchanged to its postcondition. For the lemma to apply, that piece has to be described precisely enough by F that R' is invariant under swapping it for any other satisfying F (third premise). Precision of P is required to force the same sub-state of R to satisfy both $F * \mathsf{true}$ and $Q \mathrel{-\!\!*} R'$ (second and last premises), and $R \vdash P * F * \mathsf{true}$ makes sure that R always contains $P * F$.

Finally, Lem. 4.5 and 4.6 allow one to split a ramification over either disjoint or overlapping pieces of states into independant ramifications over each of these states. This is crucial to make the proof of some ramification entailments modular, for instance when reasoning about the effects of an update on both a global graph and a set of overlapping pointers as in our proof of Cheney's algorithm (see §7), or when proving the dag copying program of §A.

This ramification library is by no means exhaustive, nor do we use all of it in the examples presented here. Rather, we think that these lemmas demonstrate that ramification entailments can be reasoned about using the intuition that they represent *updates* to the state. The lemmas we will present in §4.3 for ramifications on graphs and dags further reinforce that claim.

4.2 Exact Graph and Dag Predicates

In this section, we define mathematical graphs γ and dags δ; we will provide ways to reason about ramifications which involve them in the next section. Before we do so, however, let us consider whether our job would be any easier if we were only worried about shape instead of functional correctness, *e.g.* if we tried to verify `mark` with the spec $\{\mathtt{dag(x)}\}\ \mathtt{mark(x)}\ \{\mathtt{dag(x)}\}$.

As in Fig. 1, the proof is straightforward until the first recursive call on line 8. After framing away the root pointer $\mathrm{x} \mapsto 1, \mathrm{l}, \mathrm{r}$, we apply RAMIFY, which emits the following entailment, in which P

4.1: Frame within $\circledast$ Ramification

$$\frac{\mathit{precise}(P,Q) \qquad P \circledast R \vdash P * (Q \mathrel{-\!\!*} Q \circledast R')}{(P * F) \circledast R \vdash P * (Q \mathrel{-\!\!*} (Q * F) \circledast R')}$$

4.2: Frame within $\wedge$ Ramification 1

$$\frac{\mathit{precise}(P) \qquad R \vdash P * (Q \mathrel{-\!\!*} R')}{(P * F) \wedge R \vdash P * (Q \mathrel{-\!\!*} (Q * F) \wedge R')}$$

4.3: Frame within $\wedge$ Ramification 2

$$\frac{P \mathrel{-\!\!\circledast} R \vdash Q \mathrel{-\!\!*} R'}{(P * F) \wedge R \vdash P * (Q \mathrel{-\!\!*} (Q * F) \wedge R')}$$

4.4: Exact Frame within Ramification

$$\frac{\mathit{precise}(P) \qquad R \vdash P * F * \mathsf{true} \qquad F \mathrel{-\!\!\circledast} R' \vdash F \mathrel{-\!\!*} R' \qquad R \vdash P * (Q \mathrel{-\!\!*} R')}{R \vdash P * F * (Q * F \mathrel{-\!\!*} R')}$$

4.5: Disjoint Ramification

$$\frac{R \vdash P * (P' \mathrel{-\!\!*} R') \qquad S \vdash Q * (Q' \mathrel{-\!\!*} S')}{R * S \vdash P * Q * (P' * Q' \mathrel{-\!\!*} R' * S')}$$

4.6: $\circledast$-piecewise Ramification

$$\frac{\mathit{precise}(P,P') \qquad \forall i.\, P \circledast Q_i \vdash P * (P' \mathrel{-\!\!*} P' \circledast Q_i')}{P \circledast Q_1 \circledast Q_2 \vdash P * (P' \mathrel{-\!\!*} P' \circledast Q_1' \circledast Q_2')}$$

Figure 2: Some general-purpose lemmas from our ramification library.

and Q are the pre- and postconditions from the recursive call:

$$\overbrace{\mathsf{dag(1)} \circledast \mathsf{dag(r)}}^{R} \vdash \overbrace{\mathsf{dag(1)}}^{P} * (\overbrace{\mathsf{dag(1)}}^{Q} \mathrel{-\!\!*} \overbrace{\mathsf{dag(1)} \circledast \mathsf{dag(r)}}^{R'})$$

Unfortunately, this entailment turns out to be invalid. Recall that our ramification $P * (Q \mathrel{-\!\!*} R')$ idiom represents a state update, in which P is substituted for Q in R to yield R'. Here this means substituting one $\mathsf{dag(1)}$ for another $\mathsf{dag(1)}$, which on its surface seems reasonable. The problem is that the "ramified away" state $(Q \mathrel{-\!\!*} R')$ can have dangling pointers into the "local" state P; if P is mangled too badly as it is transformed into Q then those pointers break in the recombined state R' (here $\mathsf{dag(1)} \circledast \mathsf{dag(r)}$):

In this example, the update on dag `1` has freed node s and allocated a fresh node l'' instead. Although `1` is still a dag afterwards, `r` is not, so we will not be able to prove $\mathsf{dag(1)} \circledast \mathsf{dag(r)}$.

This is not an artificial problem stemming from our approach. In fact, the failure of the ramification entailment indicates that $\{\mathsf{dag(1)}\}$ `mark(1)` $\{\mathsf{dag(1)}\}$ is too weak of an inductive specification: overly aggressive changes to the pointer structure of the left sub-dag could make the recursive call to the right sub-dag crash, and we must reflect that reality in the specification for `mark`.

There are several solutions to this problem, but for this paper choose the most powerful: proving functional correctness. In §8.3 we will discuss some other possibilities that can yield more lightweight shape proofs at the cost of some additional formalism.

Mathematical graphs We define the mathematical representation of a directed binary graph as a quadruple (V, D, L, E), where V is a finite set of vertices, D is some set of data, $L : V \to D$ is a labeling function associating each vertex v with some data d, and $E : V \to (V \uplus \{0\}) \times (V \uplus \{0\})$ associates each vertex with up to two successors. To ease the matching between a mathematical graph and its heap representation we usually take $V \subset \mathsf{Loc}$ and $D \subseteq \mathsf{Val}$.

Given a mathematical graph $\gamma = (V, D, L, E)$, we often write $x \in \gamma$ for $x \in V \uplus \{0\}$, $S \subseteq \gamma$ for $S \subseteq V \uplus \{0\}$, and $\gamma(x)$ for $(L(x), E(x).1, E(x).2)$. We define the update of γ at node v, written $[v \mapsto (d, l, r)]\gamma$, where $l, r \in V \cup \{v\} \uplus \{0\}$ and $d \in D$, as:

$$[v \mapsto (d,l,r)](V,D,L,E) \stackrel{\text{def}}{=} (V \cup \{v\}, D, [v \mapsto d]L, [v \mapsto (l,r)]V)$$

A node y is the *successor* of a node $x \in \gamma$, written $x \overset{\gamma}{\rightsquigarrow} y$, or simply $x \rightsquigarrow y$ when γ is clear from context, if either $E(x) = (y, z)$ or $E(x) = (z, y)$ for some z. A node y is *reachable* from x, written $x \overset{\gamma}{\rightsquigarrow}{}^* y$ or $x \rightsquigarrow^* y$, if (x, y) is in the reflexive transitive closure $\rightsquigarrow$. The reachability set of $x \in \gamma$, written $\mathit{reach}(\gamma, x)$, is defined as:

$$\mathit{reach}(\gamma, x) \stackrel{\text{def}}{=} \{y \mid x \overset{\gamma}{\rightsquigarrow}{}^* y\}$$

We also lift reachability to sets of vertices $S = \{v_1, \ldots, v_n\} \subseteq V$:

$$\mathit{reach}(\gamma, S) \stackrel{\text{def}}{=} \mathit{reach}(\gamma, v_1) \cup \cdots \cup \mathit{reach}(\gamma, v_n)$$

Given a graph $\gamma = (V, D, L, E)$ and a set of vertices $S \subseteq V$, it is often useful to restrict γ to those vertices *reachable from* (respectively *not* reachable from) the vertices in S, written $\gamma \downarrow S$ (and respectively $\gamma \uparrow S$). Accordingly, we define (where $f \downharpoonright S$ is the function obtained from f by restricting the domain to the set S):

$$\gamma \downarrow S \stackrel{\text{def}}{=} (V' = \mathit{reach}(\gamma, S) \quad , D, L \downharpoonright V', E \downharpoonright V')$$
$$\gamma \uparrow S \stackrel{\text{def}}{=} (V' = V \backslash (\mathit{reach}(\gamma, S)), D, L \downharpoonright V', E \downharpoonright V')$$

The quadruple $(V', D', L', E') = \gamma \uparrow S$ is not necessarily a graph, since the edge function E' may point outside of the new set of edges V'. However, $\gamma \downarrow S$ is always a graph: the subgraph of γ reachable from S. We sometimes write $\gamma \downarrow x$ and $\gamma \uparrow x$ for $\gamma \downarrow \{x\}$ and $\gamma \uparrow \{x\}$. By convention, if $S \nsubseteq V$ then $(V, D, L, E) \downarrow S$ is the empty graph.

Spatial graphs We tie a mathematical graph γ to a spatial (in-heap) graph by adding γ as a parameter to graph:

$$\mathsf{graph}(x, \gamma) \stackrel{\text{def}}{=} \begin{array}{l}(x = 0 \wedge \mathsf{emp}) \vee \exists d, l, r.\, \gamma(x) = (d, l, r) \wedge \\ \quad x \mapsto d, l, r \circledast \mathsf{graph}(l, \gamma) \circledast \mathsf{graph}(r, \gamma)\end{array}$$

Note that $\mathsf{graph}(x, \gamma)$ "owns" only the spatial representation of the portion of γ that is reachable from x; γ may contain other nodes. This is expressed by the following lemma, that "explodes" a graph into its individual nodes. We use the *iterative star* notation, defined as emp if the set that is being iterated over is empty and as follows otherwise, given a predicate $P(x)$ on x:

$$\bigstar_{x \in \{x_1, \ldots, x_n\}} P(x) \stackrel{\text{def}}{=} P(x_1) * \cdots * P(x_n)$$

Lemma 4.7 *For every graph $\gamma = (V, E)$ and node $x \in \gamma$,*

$$\mathsf{graph}(x, \gamma) \dashv\vdash \bigstar_{v \in \mathit{reach}(\gamma, x)} v \mapsto \gamma(v)$$

Generally speaking, reasoning at this level is undesirable in program proofs, and ramification crucially allows us to remain at the level of graph instead. However, this lemma is useful both as a sanity check and to prove the general lemmas about ramifying graphs given in the next section.

We likewise enrich dag with a mathematical graph δ:

$$\mathsf{dag}(x, \delta) \stackrel{\text{def}}{=} \begin{array}{l}(x = 0 \wedge \mathsf{emp}) \vee \exists d, l, r.\, \delta(x) = (d, l, r) \wedge \\ \quad x \mapsto d, l, r * (\mathsf{dag}(l, \delta) \circledast \mathsf{dag}(r, \delta))\end{array}$$

Moreover, the predicate $\mathsf{dag}(x, \delta)$ is satisfiable *if and only if* $\delta \downarrow x$ is indeed a dag, as enforced by the $*$ in the spatial predicate:

Lemma 4.8 *For every graph δ and variable x,*

$$\mathsf{dag}(x, \delta) \dashv\vdash \mathsf{graph}(x, \delta) \wedge (\delta \downarrow x \textit{ is acyclic})$$

Finally, we define the following shorthand for describing multiple sub-graphs of the same graph from a root set $S = \{v_1, \ldots, v_n\}$:

$$\mathsf{graphs}(S, \gamma) \stackrel{\text{def}}{=} \mathsf{graph}(v_1, \gamma) \circledast \cdots \circledast \mathsf{graph}(v_n, \gamma)$$
$$\mathsf{dags}(S, \delta) \stackrel{\text{def}}{=} \mathsf{dag}(v_1, \delta) \circledast \cdots \circledast \mathsf{dag}(v_n, \delta)$$

If $S = \varnothing$ then both predicates denote emp.

4.3 Reasoning about Graph and Dag Ramifications

One advantage of proving functional correctness is that we can tightly connect our mathematical reasoning with our spatial reasoning. Here we state lemmas that do just that.

First, the spatial graph (and thus dag) predicates are precise.

Lemma 4.9 *For all S and γ, precise*$(\mathsf{graphs}(S,\gamma))$.

Our next lemma lets us *reroot* collections of sub-graphs provided that we preserve the set of reachable nodes:

Lemma 4.10 *If $reach(\gamma,S) = reach(\gamma,S')$, then*

$$\mathsf{graphs}(S,\gamma) \dashv\vdash \mathsf{graphs}(S',\gamma)$$

Our third lemma helps us extend a graph with fresh nodes.

Lemma 4.11 (Graph Growth)

$$x \mapsto d,x,x \vdash \mathsf{graph}(x,[x\mapsto(d,x,x)]) \quad (13)$$

$$x \mapsto d,x,r * \mathsf{graph}(r,\gamma) \vdash \mathsf{graph}(x,[x\mapsto(d,x,r)]\gamma) \quad (14)$$

$$x \mapsto d,l,r * \mathsf{graphs}(\{l,r\},\gamma) \vdash \mathsf{graph}(x,[x\mapsto(d,l,r)]\gamma) \quad (15)$$

$$x \mapsto d,l,r * \mathsf{dags}(\{l,r\},\delta) \vdash \mathsf{dag}(x,[x\mapsto(d,l,r)]\delta) \quad (16)$$

First, (13) a graph cell x whose successors are both itself corresponds to a singleton graph $[x \mapsto (d,x,x)]$.[4] Second, (14) if a node x has a loop to itself on the left and a pointer to an existing graph on the right,[5] then we can add x to the graph; not shown is the mirrored case when the loop is on the right. Third, (15) if x links to two (possibly equal) graph nodes then we can again add x to the graph. The first two cases need to be stated separately because $x \notin \gamma$ means that $\mathsf{graph}(x,\gamma)$ is $\bot$. Finally, (16) is the analogue of (15) for dags; we do not need analogues for (13) and (14) because dags must be acyclic.

The frame rule, combined with the $*$ between a parent and its descendants and equation (16), is enough to mutate the root of a dag. However, an unrestricted graph has $\circledast$ between the parent and its successors, so we need to use RAMIFY to update the root. The following lemma helps discharge the associated ramifications:

Lemma 4.12 (Single Graph Node Update)

$$\frac{\gamma(x) = (d,l,r) \qquad \gamma' = [x\mapsto(d',l',r')]\gamma}{\mathsf{graphs}(\{x,l',r'\}\cup S,\gamma) \vdash x\mapsto d,l,r * (x\mapsto d',l',r' \mathrel{-\!\!*} \mathsf{graphs}(\{x,l,r\}\cup S,\gamma'))} \quad (17)$$

$$\frac{\gamma(x) = (d,l,r) \qquad \gamma' = [x\mapsto(d',l,r)]\gamma}{\mathsf{graph}(x,\gamma) \vdash x\mapsto d,l,r * (x\mapsto d',l,r \mathrel{-\!\!*} \mathsf{graph}(x,\gamma'))} \quad (18)$$

Lem. 4.11 handles the cases in which we are adding a fresh node, so in (17)-(18) we need only consider the case in which $\{x,l,r,l',r'\} \subseteq \gamma$. The case of interest is (17), a full update to node x, where we are updating not only the data d to d' but also the pointers l and r to l' and r' respectively. The precondition is a $\circledast$-joined set of subgraphs of γ, including x, l', and r', as well as arbitrary others S. After the update, the state contains subgraphs at x (which now contains l' and r') and S, as well as the old l and r (previously contained in the old x), which may now be disconnected from $\{x\} \cup S$. In practice we often care about far simpler updates; (18) is a direct consequence of (17), and handles the case in which we only wish to update the data field.

Next, we observe that an update that preserves the set of reachable nodes cannot remove any overlapping points-to fact. The same remark is true of dags as well, replacing graphs with dags everywhere in the lemma below.

[4] $[x\mapsto(d,x,x)] \stackrel{\text{def}}{=} (\{x\}, D, [x\mapsto d], [x\mapsto(x,x)])$

[5] Observe that r can be equal to 0, in which case $\mathsf{graph}(r,\gamma)$ is just emp.

Lemma 4.13 (Points-to preservation)

$$\frac{reach(\gamma',S') \supseteq reach(\gamma,S)}{\mathsf{graphs}(S,\gamma) \circledast x\mapsto - \vdash \mathsf{graphs}(S,\gamma) * (\mathsf{graphs}(S',\gamma') \mathrel{-\!\!*} \mathsf{graphs}(S',\gamma') \circledast x\mapsto -)}$$

Our final lemma applies when we wish to update an entire subgraph (typically with a function call) rather than a single node.

Lemma 4.14 (Subgraph Update)

$$\frac{reach(\gamma',S'_1) \supseteq reach(\gamma,S_1) \qquad \gamma' \uparrow S'_1 = \gamma \uparrow S_1}{\mathsf{graphs}(S_1,\gamma) \circledast \mathsf{graphs}(S_2,\gamma) \vdash \mathsf{graphs}(S_1,\gamma) * (\mathsf{graphs}(S'_1,\gamma') \mathrel{-\!\!*} \mathsf{graphs}(S'_1,\gamma') \circledast \mathsf{graphs}(S_2,\gamma'))} \quad (19)$$

$$\frac{reach(\delta',S'_1) \supseteq reach(\delta,S_1) \qquad \delta' \uparrow S'_1 = \delta \uparrow S_1}{\mathsf{dags}(S_1,\delta) \circledast \mathsf{dags}(S_2,\delta) \vdash \mathsf{dags}(S_1,\delta) * (\mathsf{dags}(S'_1,\delta') \mathrel{-\!\!*} \mathsf{dags}(S'_1,\delta') \circledast \mathsf{dags}(S_2,\delta'))} \quad (20)$$

First, (19) lets us ramify an update to a subgraph (or set of subgraphs) as long as all previously reachable nodes are still reachable (to prevent *e.g.* the dangling pointer problem outlined in §3.3) and the mathematical update is local. Second, (20) gives us the same property for dags (if our newly substituted sub-dag does not contain a cycle then our whole dag will not suddenly become cyclic).

5. Proving `mark` on Dags

We are ready at last to polish off the proof of `mark` from Fig. 1.

Mathematical marking One of our goals is to translate mathematical reasoning into spatial reasoning. Define the mathematical marking $m(\gamma,r)$ of a graph $\gamma = (V,D,L,E)$ starting from the vertex $r \in V$ as marking all nodes reachable via *unmarked nodes* from r. Formally, define a new relation $\overset{\gamma}{\rightsquigarrow}_0$ as follows:

$$x \overset{\gamma}{\rightsquigarrow}_0 y \text{ iff } \exists z.\, \gamma(x) = (0,y,z) \vee \gamma(x) = (0,z,y)$$

As before, we omit the subscript γ when it is clear from context and write $\overset{\gamma}{\rightsquigarrow}_0^*$ for the reflexive transitive closure. The marking $m(\gamma,r)$ of γ from r is then (V,D,L',E) where, for all $x \in V$,

$$L'(x) \;=\; \begin{cases} 1 & \text{if } r \rightsquigarrow_0^* x \\ L(x) & \text{otherwise} \end{cases}$$

We also need to describe the effect of marking a single node in γ, accomplished with $m_1(\gamma,x)$, that sets the marked bit of node x in γ to 1. The following lemma about mathematical markings now becomes crucial to prove the functional correctness of `mark`.

Lemma 5.1 *For all graphs γ and nodes $x,y \in \gamma$,*

$$m(m(\gamma,x),y) \;=\; m(m(\gamma,y),x) \quad (21)$$

Moreover, if $\gamma(x) = (d,l,r)$, then

$$m(m(m_1(\gamma,x),l),r) = m(m_1(m(\gamma,l),x),r) = m_1(m(m(\gamma,l),r),x) = m(m_1(m(\gamma,r),x),l) = m(\gamma,x) \quad (22)$$

That is, (21) we can swap the order of two mathematical markings, and (22) regardless of which order we mark the root and children (either child first by equation 21), at the end we are fully marked.

Spatial marking Our first remaining tasks are the ramifications on lines 9 and 11. In both cases we frame away the root node and then apply RAMIFY, yielding the following entailments:

$$\mathsf{dag}(\mathtt{l},\delta) \circledast \mathsf{dag}(\mathtt{r},\delta) \vdash \mathsf{dag}(\mathtt{l},\delta) * (\mathsf{dag}(\mathtt{l},m(\delta,\mathtt{l})) \mathrel{-\!\!*} \mathsf{dag}(\mathtt{l},m(\delta,\mathtt{l})) \circledast \mathsf{dag}(\mathtt{r},m(\delta,\mathtt{l}))) \quad (23)$$

$$\mathsf{dag}(\mathtt{l},m(\delta,\mathtt{l})) \circledast \mathsf{dag}(\mathtt{r},m(\delta,\mathtt{l})) \vdash \mathsf{dag}(\mathtt{r},m(\delta,\mathtt{l})) * (\mathsf{dag}(\mathtt{r},m(m(\delta,\mathtt{l}),\mathtt{r})) \mathrel{-\!\!*} \mathsf{dag}(\mathtt{l},m(m(\delta,\mathtt{l}),\mathtt{r})) \circledast \mathsf{dag}(\mathtt{r},m(m(\delta,\mathtt{l}),\mathtt{r}))) \quad (24)$$

Observe that the first ramification directly implies the second by instantiating δ with $m(\delta,\mathtt{l})$ in the first entailment and using the

commutativity of $\circledast$ to swap the roles of `l` and `r`. To prove (23), we apply Lem. 4.14 to reduce the spatial ramification entailment to a pair of mathematical subgoals. The first mathematical subgoal,

$$reach(m(\delta, \mathtt{l}), \mathtt{l}) \supseteq reach(\delta, \mathtt{l}),$$

i.e., that every vertex reachable from `l` in the old dag δ is still reachable from `l` in the new dag $m(\delta, \mathtt{l})$, is immediate because the mathematical marking function m changes neither the vertices nor the edges of the dag δ. The second mathematical subgoal,

$$m(\delta, \mathtt{l}) \upharpoonright \mathtt{l} = \delta \upharpoonright \mathtt{l},$$

i.e., that the part of δ that is *not* reachable from `l` is identical to the part of $m(\delta, \mathtt{l})$ that is *not* reachable `l`, is almost as simple. By the definition of m, we know that the only difference between δ and $m(\delta, \mathtt{l})$ is that the new labeling function has marked vertices reachable via unmarked paths in δ from `l`; all other labels are maintained. Since the second mathematical subgoal only cares about changes to the portion of the mathematical graph that is not reachable from `l`, and those labels are unchanged, we are done.

Finally, to establish the postcondition in line 12 from line 11, apply Lem. 4.11 to derive $\mathsf{dag}(\mathtt{x}, m_1(m(m(\delta, \mathtt{l}), \mathtt{r}), \mathtt{x}))$, which by Lem. 5.1 is equivalent to our postcondition.

5.1 Observations

Our proof of `mark` (*i.e.*, Fig. 1 and §5) is short and our invariants at each program point are straightforward. We were able to reuse our initial ramification (23) to prove our second (24). Essentially all of the spatial difficulties were handled by our ramification library. Moreover, by Lem. 5.1 our proof is easy to modify to accommodate trivial changes in the program like moving the update in line 6 to after one or both of the recursive calls in lines 8 and 10, swapping the order of the recursive calls, etc. Our ability to accommodate these kinds of changes is an indication of the power of using ramification to separate mathematical and spatial reasoning from each other.

In contrast, previous work on verifying these kinds of algorithms used complex and brittle invariants so that they could always apply the frame rule. For example, consider Bornat *et al.* [BCO04], which is the progenitor of most previous work in applying separation logic to reason about data structures with intrinsic sharing.

Bornat *et al.* define mathematical dags as tree-shaped terms whose nodes are either labeled proper nodes (written $x : \mathsf{Node}\ \delta_l\ \delta_r$) or references to a label elsewhere in the dag (written $\mathsf{Ptr}\ x$) to express sharing. Their spatial pdag predicate grants ownership of a node at the point that corresponds to where it is declared in the mathematical definition (roughly speaking, $\mathsf{pdag}\ x\ (x : \mathsf{Node}\ \delta_l\ \delta_r) \stackrel{\text{def}}{=} x \mapsto v, l, r * \mathsf{pdag}\ l\ \delta_l * \mathsf{pdag}\ r\ \delta_r$), but not when it is referenced again (similarly, $\mathsf{pdag}\ x\ (\mathsf{Ptr}\ x) \stackrel{\text{def}}{=} \mathsf{emp}$). Each node can only be declared once (although it may be referenced many times), and the order in which they are declared must match the order in which the program traverses the dag, as the authors note [BCO04, §8, p. 7]:

> This predicate is specifically designed to support a left to right scan, as are the formulae on which it is based. It seems difficult to avoid this complication.

Unfortunately, the consequences of this style of definition ricochet to many other parts of the associated verification, including the statement of the specification of `mark` and its exact implementation. Changes in one part of the system (*e.g.*, swapping the order of traversal) required changes to other parts of the system (*i.e.*, the definition of pdag, the specification of `mark`, and the invariants at each program point). Altogether, the style of hacking the state into many disjoint pieces to reason about data structures with intrinsic sharing pays a heavy price to enable the frame rule, resulting in "dauntingly subtle" [BCO04, §8.4, p. 9] definitions and verifications.

In contrast, our definition of mathematical graphs is traditional, our dag predicate is natural, and our specification for `mark` is straightforward. None of these depend on internal implementation specifics of the algorithm such as traversal orders. Moreover, our program invariants are easy to understand, easy to update to accommodate minor changes in the algorithm, and easy to verify using ramification and our ramification library. Verifications utilizing ramifications are both more natural and more robust than those in previous work.

Marking possibly cyclic graphs The `mark` function can also mark unrestricted graphs. Because Lemmas 4.14 and 5.1 both apply to graphs as well as dags, the only substantial change to the the proof in Fig. 1 is for line 6. Here dags only require the frame rule due to the $*$ between a parent and its children but unrestricted graphs require an additional ramification due to the additional $\circledast$:

$$\mathtt{x} \mapsto 0, \mathtt{l}, \mathtt{r} \circledast \mathsf{graphs}(\{\mathtt{l}, \mathtt{r}\}, \gamma) \vdash \mathtt{x} \mapsto 0, \mathtt{l}, \mathtt{r} * (\mathtt{x} \mapsto 1, \mathtt{l}, \mathtt{r} \mathrel{-\!\!*} \mathtt{x} \mapsto 1, \mathtt{l}, \mathtt{r} \circledast \mathsf{graphs}(\{\mathtt{l}, \mathtt{r}\}, m_1(\gamma, \mathtt{x})))$$

This ramification follows directly from Lem. 4.12.

Termination Our work here is primarily concerned with partial correctness, but suppose we were interested in total correctness as well. The dag argument is simpler: each recursive call is on a strictly smaller subheap thanks to the $*$ between a parent and its children; notice that this argument is valid regardless of whether we mark the root first, at line 6, or after one or both recursive calls. In contrast, the termination argument on unrestricted graphs is more complicated because the $\circledast$ between root and successors means that the subheap may not be any smaller at the recursive calls. Instead, each recursive call must be on a graph with fewer unmarked nodes; if we recurse before coloring the root then we may not terminate.

5.2 Other Graph Algorithms

To prove that ramification can apply equally well to programs that, unlike `mark`, mutate the link structure of the graph, we also verified `copy_dag` and `dispose_graph`. The full details are in Appendices A and B respectively; here we give only the key insights.

Copying dags The goal of the `copy_dag` function is to make a deep (structure-preserving) copy of its argument. It uses a data field in each original node to record the location of its corresponding copy (or 0 if the node has not yet been copied). Just as with `mark`, a straightforward recursive implementation is compact and works as follows. If the root is already copied then return immediately; otherwise, recursively copy the left and right children, allocate a new node to be the root's copy, and set its fields as appropriate.

To make the verification hang together, we need to add a new feature to our separation logic: *regions* [LG88]. Briefly, regions indicate disjoint zones in the heap, and a spatial predicate P can be tagged with a region identifier α to become P_α, indicating that P is entirely contained in region α. Predicates in different regions are always disjoint, even when connected by sharing operators such as $\circledast$.

Regions are useful when we are faced with the following problem, in which $\sharp$ is some sharing operator such as $\wedge$ or $\circledast$:

$$(P * Q) \sharp (R * S) \ \dashv\stackrel{?}{\vdash}\ (P \sharp R) * (Q \sharp S)$$

That is, we have some disjoint formulas P and Q, which overlap with two additional disjoint formulas R and S, and we wish to shuffle resources around until the P and Q are overlapping with each other and are disjoint from the overlapping Q and S.

The $\dashv$ direction is immediate. Unfortunately, verifying `copy_dag` requires the $\vdash$ direction after making both recursive calls and reaching the following invariant, in which `l` and `r` are the left and right children of the root and `ll` and `rr` are their respective copies:

$$(\mathsf{dag}(\mathtt{l}, \delta) * \mathsf{dag}(\mathtt{ll}, \delta')) \circledast (\mathsf{dag}(\mathtt{r}, \delta) * \mathsf{dag}(\mathtt{rr}, \delta'))$$

Now we need to apply the rule of consequence to disentangle the original children from their overlapping copies:

$$(\mathsf{dag}(\mathtt{l}, \delta) \circledast \mathsf{dag}(\mathtt{r}, \delta)) * (\mathsf{dag}(\mathtt{ll}, \delta') \circledast \mathsf{dag}(\mathtt{rr}, \delta'))$$

```
1  struct node { struct node *next,*l,*r; };
2  void pop(void) { //{list(s) ∧ tree(t)}
3    if (!s) return;
4    struct node *c = s;
5  // {(s ↦ n, l, r ∗ list(n)) ∧ tree(t) ∧ c = s}
6    s = c->next;
7  // {(c ↦ s, l, r ∗ list(s)) ∧ tree(t)}
8  // {(c ↦ s, l, r ∗ list(s)) ∧ (sktree(t, π ⊎ {c}) ∗ ptrs(π ⊎ {c}))}
9    t = tree_del(t,c);
10 //↯{(c ↦ s, −, − ∗ list(s)) ∧ (sktree(t, π) ∗ c ↦ −, −, − ∗ ptrs(π))}
11 // {(list(s) ∧ tree(t)) ∗ c ↦ −, −, −}
12   free(c);
13 } // {list(s)) ∧ tree(t)}
```

Figure 3: Removal from a threaded tree.

The problem is that the implication is just not true without carrying around some additional information via regions: specifically, that the original dag is in region α while the copy is being created in region β. In general regions help because they ensure that P does not have any overlap with S despite the intermediate sharing operator $\sharp$:

$$(P_\alpha * Q_\beta)\sharp(R_\alpha * S_\beta) \dashv\vdash (P\sharp R)_\alpha * (Q\sharp S)_\beta \qquad (25)$$

Others have run across the same problem in diverse contexts including RGSep [Vaf07] and shape analysis for overlaid lists and trees [LYP11] and have turned to regions for similar reasons.

Interestingly, our verification also uses regions in an novel way to split one large ramification entailment (equation 35) into two smaller entailments via Lem. 4.5 from our ramification library. This second use of regions is not vital to verify `copy_dag`, but it does simplify things nicely. Other than the use of regions, the verification proceeds straightforwardly.

Disposing graphs Disposing a graph is usually done in two steps: first, suppress all sharing between nodes of the graph, so that each node has at most one predecessor, thus computing a spanning tree of the graph; and then dispose the tree. Apx. B contains the novel verification for the first step; verifying the second is standard. The real proof effort is on the mathematical side; the spatial aspects of the verification are no more complicated than `mark` and `copy_dag`, and do not require regions. Because our definition for graph uses $\circledast$, we are able to establish emp at the end, indicating that we have completely freed the structure.

6. Overlaid Data Structures

Reasoning about threaded trees Our examples so far have focused on graph manipulations. Ramification is also applicable in other interesting contexts, including overlaid data structures. Here we focus on one kind of overlaid structure: *threaded trees*, which overlay lists and trees. Each node has *three* links to other nodes of the data structure: a "next" pointer of a singly-linked list, and the "left" and "right" fields of a binary tree. This is a popular type of overlaid data structure: the linked list may record the set of elements some order of particular interest (*e.g.*, first-inserted to most recent), while the tree provides efficient out-of-order lookup.

Our case study is a procedure that removes the first element of the linked list from the data structure, inspired by what can be found in the Linux deadline I/O scheduler [LYP11]. The code and annotations are shown in Fig. 3. It assumes two global variables s and t that point respectively to the head of the linked list and the root of the tree. The precondition states that the two shapes span *exactly* the same memory cells, enforced by the conjunction $\wedge$. Removing from the list (line 6) merely advances the head pointer, but we cannot stop there because it leaves the overlaid structure in an inconsistent state (the items in the list and the tree must be identical).

Removing from the tree is likely to be operationally complex, potentially involving operations to rebalance, reroot, or otherwise rotate parts of the tree. Thus, we abstract this operation and assume that it is performed by a function `tree_del(t,c)`. Its spec has to express two particular facts to ensure that it is well-behaved w.r.t. the overlaid list structure: it must not tamper with the list fields, and the resulting new tree should cover the same nodes as before except for `c`. We enforce the first constraint by not giving any access rights on the list fields to the procedure, *i.e.* by restricting its precondition to the "skeleton" of the tree, and the second constraint by recording the set of nodes encompassed in the tree shape. We therefore define the following predicate that skips the list fields of each node:

$$\begin{aligned}\mathsf{sktree}(x,\pi) \stackrel{\text{def}}{=} {} & (x = 0 \wedge \mathsf{emp} \wedge \pi = \varnothing) \vee \exists l, r, \pi_l, \pi_r.\\ & x+1 \mapsto l, r * \mathsf{sktree}(l, \pi_l) * \mathsf{sktree}(r, \pi_r) \wedge\\ & \pi = \{x\} \uplus \pi_l \uplus \pi_r\end{aligned}$$

The tree predicate can be split into a skeleton and a bag of points-to predicates, using the *pointers* predicate ptrs:

$$\begin{aligned}\mathsf{ptrs}(\{x_1,\ldots,x_n\}) \stackrel{\text{def}}{=} {} & x_1 \mapsto - * \cdots * x_n \mapsto -\\ \mathsf{tree}(t) \Leftrightarrow {} & \exists \pi.\, \mathsf{sktree}(t,\pi) * \mathsf{ptrs}(\pi)\end{aligned} \qquad (26)$$

The list predicate is defined in the standard way for nil-terminated acyclic lists with two data fields:

$$\mathsf{list}(l) \stackrel{\text{def}}{=} (l = 0 \wedge \mathsf{emp}) \vee \exists l', x, y.\, l \mapsto l', x, y * \mathsf{list}(l')$$

We moreover assume that each address is aligned as a multiple of 3, to prevent *skewing*, in which a node in the tree might overlap two nodes in the list in a state satisfying $\mathsf{list}(s) \wedge \mathsf{tree}(t)$.

A general observation about how overlaid data structures are manipulated is that changes to fields of only one structure do not affect the other, *e.g.*, list induction easily proves that

$$x \mapsto n, l, r \mathrel{-\!\!\circledast} \mathsf{list}(s) \vdash x \mapsto n, l', r' \mathrel{-\!\!*} \mathsf{list}(s)$$

This reads as: if a state may be completed by a node to form a linked list, then completing it by any other node at the same location and with the same next field also yields a list. The same property for *skeleton* trees follows by induction on the size of the tree:

$$\mathsf{sktree}(t,\pi) \mathrel{-\!\!\circledast} \mathsf{list}(s) \vdash \mathsf{sktree}(t',\pi) \mathrel{-\!\!*} \mathsf{list}(s) \qquad (27)$$

Verification The spec of `tree_del` follows the discussion above:

$$\{\mathsf{sktree}(t, \pi \uplus \{c\})\}\texttt{u=tree\_del(t,c)}\{\mathsf{sktree}(u,\pi) * c{+}1 \mapsto -, -\}$$

The proof sketched in Fig. 3 is mostly straightforward: if `s` is nil then the list is empty, hence so is the tree and the postcondition is trivially satisfied; otherwise, we unfold the list predicate, which enables the lookup at line 6. After that, we split the tree according to (26) and apply the following ramification:

$$\begin{aligned}&(\mathtt{c} \mapsto \mathtt{s}, l, r * \mathsf{list}(\mathtt{s})) \wedge (\mathsf{sktree}(\mathtt{t}, \pi \uplus \{\mathtt{c}\}) * \mathsf{ptrs}(\pi \uplus \{\mathtt{c}\}))\\ \vdash {} & \mathsf{sktree}(\mathtt{t}, \pi \uplus \{\mathtt{c}\}) * \big(\mathtt{c}+1 \mapsto -, - * \mathsf{sktree}(t', \pi) \mathrel{-\!\!*}\\ & \quad (\mathtt{c} \mapsto \mathtt{s}, -, - * \mathsf{list}(\mathtt{s})) \wedge\\ & \quad (\mathtt{c}+1 \mapsto -, - * \mathsf{sktree}(t', \pi) * \mathsf{ptrs}(\pi \uplus \{\mathtt{c}\}))\big)\end{aligned}$$

This ramification follows a general pattern, and we can reduce it to a much simpler one by noticing that the right-hand side conjunct is automatically handled by Lem. 4.3 from our ramification library, which can remove frames that occur within $\wedge$ ramifications. This yields the following simpler proof obligation:

$$\begin{aligned}&\mathsf{sktree}(t, \pi \uplus \{c\}) \mathrel{-\!\!\circledast} c \mapsto x, y, z * \mathsf{list}(s)\\ \vdash {} & c+1 \mapsto -, - * \mathsf{sktree}(t', \pi) \mathrel{-\!\!*} c \mapsto x, -, - * \mathsf{list}(s)\end{aligned}$$

This entailment is similar to (27). The rest of the proof is immediate.

7. Cheney's Garbage Collector

It is time for the acid test: verifying the functional correctness of Cheney's garbage collector [Che70]. The general setting is as follows. There are two disjoint, equally large regions of memory, the *from-space* and the *to-space*, starting respectively at the address pointed to by `from` and `to`. Programs manipulate *objects* in the from-space. When the program wishes to allocate but the from-space has run out of room, we garbage collect by copying the entire graph

of reachable objects into the to-space before swapping `from` and `to` and resuming normal execution. If the former from-space had any unreachable objects then the new from-space has some free space.

In the tradition of previous work, we make a number of simplifications. We assume that there is a single root from which all *active* objects are reachable, *i.e.* any object *not* reachable from that root can be safely reclaimed. We also restrict our study to even-aligned two-field objects that contain only pointers (including the null pointer) rather than arbitrary integers. Our proof can be modified to verify the unsimplified algorithm; *e.g.*, we can allow data by the usual systems hack of requiring that data be odd and pointers be even.

Remarkably, Cheney's algorithm migrates the graph from one space to the other using only a *constant* amount of extra memory, which is in short supply during garbage collection. Contrast this with our dag-copying example of §A that required *linear* additional space (in both the data fields and the function stack). The cost is that we mangle the original graph, which we can live with because afterwards it will be garbage. The trick is that Cheney rewires the first field in each already-copied object in the from-space to point to its copy in the to-space. The collector can determine whether an object has already been copied, and moreover discover the copy's address, by checking if its first field points into the to-space.

Following [Gas11], we implement the algorithm as two functions: `collect` and `copy_ref`, shown in Fig. 4. In addition to the `to` and `from` pointers (fixed for the duration of the collection), they maintain two additional pointers into the to-space. First, the `scan` pointer separates the fully-processed "scanned" objects, whose pointers point into the to-space, from the partially-processed "queued" objects, whose pointers point into the from-space. Second, the `free` pointer distinguishes the first unused address in the to-space.

Initially (line 3), `scan = free = to`, meaning that no objects have been copied and the entire to-space is free. The process is initiated by copying the object pointed to by the root `r` (line 4), which allocates two cells of memory at the beginning of the to-space by increasing `free` and fills them with the values in the original object, now enqueued. After that the program loops (lines 5–12) until no queued objects remain, calling `copy_ref` on both object fields (lines 8 and 11) before incrementing `scan` to indicate that the object has been scanned. Each call to `copy_ref` swings the from-space pointers into the to-space, queuing newly encountered nodes as necessary. Fig. 5 presents an intermediate state in the execution, with one node copied and scanned and one node queued for scanning.

Formal specification To represent states of the execution we use the following definitions. Mathematical graphs are pairs (V, E), *i.e.* we remove D and L, and the spatial predicate is accordingly

$$\mathsf{graph}(x,\gamma) \stackrel{\text{def}}{=} \begin{array}{l}(x = 0 \wedge \mathsf{emp}) \vee \exists l, r.\, \gamma(x) = (l, r) \wedge \\ x \mapsto l, r \mathbin{\text{⍟}} \mathsf{graph}(l,\gamma) \mathbin{\text{⍟}} \mathsf{graph}(r,\gamma)\end{array}$$

We define shorthand to express whether a node is in the from- or to-space and whether it has been copied (recall that `to`, `from` and `size` are constant throughout the execution):

$$\begin{aligned} \mathit{from}(x) &\stackrel{\text{def}}{=} x = 0 \vee \mathtt{from} \leqslant x < \mathtt{from} + \mathtt{size} \\ \mathit{to}(x) &\stackrel{\text{def}}{=} x = 0 \vee \mathtt{to} \leqslant x < \mathtt{to} + \mathtt{size} \\ \mathit{copied}(\gamma, x) &\stackrel{\text{def}}{=} x \neq 0 \wedge \mathit{from}(x) \wedge \mathit{to}(\gamma(x).1) \end{aligned}$$

We write $\mathit{from}(\gamma)$ for $\forall v \in \gamma.\, \mathit{from}(v)$ and similarly for $\mathit{to}(\gamma)$. The memory also contains a *pool* of free addresses, starting at some x, and the whole from-space, which we use to collect nodes that the algorithm disconnects (*i.e.*, from-space objects that are no longer reachable from the to-space and are therefore fresh garbage):

$$\begin{aligned} \mathsf{pool}(x) &\stackrel{\text{def}}{=} \mathsf{ptrs}(\{x, \ldots, \mathtt{to} + \mathtt{size} - 1\}) \\ \mathsf{fromsp} &\stackrel{\text{def}}{=} \mathsf{ptrs}(\{\mathtt{from}, \ldots, \mathtt{from} + \mathtt{size} - 1\}) \end{aligned}$$

The main end-to-end property of a garbage collector is that the final graph is *isomorphic* to the original one. In the middle of a

```
1  void collect(void **r) {
2  // {(r ↦ r0 ⍟ graph(r0,γ0) ⍟ fromsp) ⍟ pool(to) ∧ from(γ0)}
3    scan = free = to;
4    copy_ref(r);
5    while (scan != free)
6  // { r ↦ to ⍟ (graph(to,γ) ⍟ fromsp) ⍟ pool(free) ∧
   //   γ@to ≈ γ0@r0 ∧ cheney(γ, scan, free) }
7  { // { r ↦ to ⍟ (graph(to,γ) ⍟ scan ↦ q0,q1 ⍟ graph(q0,γ) ⍟
   //     graph(q1,γ) ⍟ fromsp) ⍟ pool(free) ∧ γ@to ≈ γ0@r0 ∧
   //     cheney(γ, scan, free) ∧ scan ⩽ free − 2 }
8      copy_ref(scan);
9  //↯ { r ↦ to ⍟ (graph(to,γ′) ⍟ scan ↦ q0′,q1 ⍟ graph(q0′,γ′) ⍟
   //     graph(q0,γ′) ⍟ graph(q1,γ′) ⍟ fromsp) ⍟ pool(free) ∧
   //     γ@to ≈ γ0@r0 ∧ scan ⩽ free − 2 ∧
   //     cheney(γ′, scan + 1, free) ∧ γ′@to ≈ γ@to }
10 // { r ↦ to ⍟ (graph(to,γ′) ⍟ fromsp) ⍟ pool(free) ∧
   //   γ′@to ≈ γ0@r0 ∧ cheney(γ′, scan + 1, free) ∧
   //   scan ⩽ free − 2 }
11     copy_ref(scan + 1);
12     scan = scan + 2; }
13 } // { r ↦ to ⍟ graph(to,γ) ⍟ fromsp ⍟ pool(free) ∧
   //   to(γ) ∧ γ@to ≈ γ0@r0 }
14
15 void copy_ref(void **p) {
16 // {(p ↦ q ⍟ graph(q,γ)) ⍟ pool(f) ∧ cheney(γ, p, f) ∧ free = f}
17   if (*p) {
18 // { (p ↦ q ⍟ q ↦ a,b ⍟ graph(a,γ) ⍟ graph(b,γ)) ⍟
   //   pool(f) ∧ cheney(γ, p, f) ∧ free = f }
19     void *obj = *p;
20     void *fwd = *obj;
21 // { (p ↦ obj ⍟ obj ↦ fwd,b ⍟ graph(fwd,γ) ⍟ graph(b,γ)) ⍟
   //   pool(f) ∧ cheney(γ, p, f) ∧ free = f }
22     if (to <= fwd && fwd < to + size){
23 // { (p ↦ obj ⍟ obj ↦ fwd,b ⍟ graph(fwd,γ) ⍟ graph(b,γ)) ⍟
   //   pool(f) ∧ cheney(γ, p, f) ∧ free = f ∧ to(fwd) }
24       *p = fwd;
25 //↯ { (p ↦ fwd ⍟ obj ↦ fwd,b ⍟ graph(fwd,γ′) ⍟ graph(b,γ′)) ⍟
   //     pool(f) ∧ cheney(γ, p, f) ∧ free = f ∧ to(fwd) ∧
   //     cheney(γ′, p + 1, f) ∧ γ′ = [p ↦ fwd]γ }
26     } else {
27 // { (p ↦ obj ⍟ obj ↦ fwd,b ⍟ graph(fwd,γ) ⍟ graph(b,γ)) ⍟
   //   pool(f) ∧ cheney(γ, p, f) ∧ free = f ∧ from(fwd) }
28       void *new = free;
29       free = free + 2;
30       *new = *obj;
31       *(new + 1) = *(obj + 1);
32 // { (p ↦ obj ⍟ obj ↦ fwd,b ⍟ graph(fwd,γ) ⍟ graph(b,γ)) ⍟
   //   new ↦ fwd,b ⍟ pool(free) ∧ free = f + 2 ∧
   //   cheney(γ, p, f) ∧ from(fwd) ∧ new = f }
33 // { (p ↦ obj ⍟ obj ↦ fwd,b ⍟ new ↦ fwd,b ⍟
   //   graph(fwd,γ1) ⍟ graph(b,γ1)) ⍟ pool(free) ∧
   //   free = f + 2 ∧ cheney(γ, p, f) ∧ from(fwd) ∧
   //   new = f ∧ γ1 = [new ↦ fwd,b]γ }
34       *obj = new;
35 //↯ { (p ↦ obj ⍟ obj ↦ new,b ⍟ new ↦ fwd,b ⍟
   //     graph(fwd,γ2) ⍟ graph(b,γ2)) ⍟ pool(free) ∧
   //     cheney(γ, p, f) ∧ free = f + 2 ∧ from(fwd) ∧
   //     new = f ∧ γ2 = [obj ↦ new][new ↦ fwd,b]γ }
36       *p = new;
37 //↯ { (p ↦ new ⍟ obj ↦ new,b ⍟ new ↦ fwd,b ⍟
   //     graph(fwd,γ′) ⍟ graph(b,γ′)) ⍟ pool(free) ∧
   //     cheney(γ, p, f) ∧ free = f + 2 ∧ from(fwd) ∧
   //     new = f ∧ γ′ = [p ↦ new][obj ↦ new][new ↦ fwd,b]γ }
38 }}} // { (p ↦ q′ ⍟ graph(q′,γ′) ⍟ graph(q,γ′)) ⍟ pool(free) ∧
   //   cheney(γ′, p + 1, free) ∧ γ@to ≈ γ′@to ∧
   //   γ′ ↑ {p, q, q′} = γ ↑ {p, q} ∧ free ⩾ f }
```

Figure 4: Proof sketch of Cheney's garbage collector.

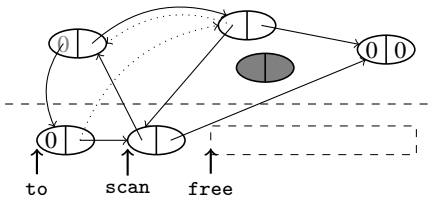

Figure 5: Transient state of the memory during garbage collection. Previous field values are indicated by 0 or dotted pointer arrows.

$$cheney(\gamma, s, f) \stackrel{\text{def}}{=} to(s) \wedge to(f) \wedge$$

$$|\{v \mid copied(\gamma, v)\}| = (f - \texttt{to})/2 \wedge \tag{28}$$

$$\{\texttt{to}, \ldots, f-2\} \subseteq \gamma \downarrow \texttt{to} \wedge \tag{29}$$

$$\forall v \in \gamma.\, \forall a, b.\, \gamma(v) = (a, b) \Rightarrow$$

$$\begin{pmatrix} to(v) \wedge ((v < s \wedge to(a)) \vee (v \geqslant s \wedge from(a))) \wedge \\ ((v{+}1 < s \wedge to(b)) \vee (v{+}1 \geqslant s \wedge from(b))) \end{pmatrix} \vee \tag{30}$$

$$\big(from(v) \wedge from(b) \wedge (to(a) \Rightarrow \gamma@b \approx \gamma@(\gamma(a).2))\big) \tag{31}$$

Figure 6: Cheney graphs. Parameter s is the first unscanned address and f is the beginning of the free space. There are as many nodes in the to-space as there are copied nodes (28), which ensures that we never exhaust our free space at line 29. Every cell in the to-space is reachable (29). For each object, either (30) it is in the to-space and either scanned with fields pointing to the to-space, or queued with fields pointing to the from-space; or (31) it is in the from-space, with fields either entirely pointing to the from-space or with first field pointing to its copy in the to-space, in which case the second fields of the object and its copy point to isomorphic sub-graphs.

collection, the loop invariant is more complex; for Cheney it is that the graph rooted at `to` is isomorphic to the original one *up-to* a canonicalization function, $canon(\gamma)$. The canonicalization of a graph $\gamma = (V, E)$ "skips" already copied nodes by following their first field (which points to their copies). Formally, $canon(\gamma)$ is the graph (V', E') where $V' = \{v \in V \mid \neg copied(\gamma, v)\}$ and, if $E(x) = (v_1, v_2)$, then $E'(x) = (v_1', v_2')$ with

$$v_i' = \begin{cases} E(v_i).1 & \text{if } copied(\gamma, v_i) \\ v_i & \text{otherwise} \end{cases}$$

We write $\gamma@x \approx \gamma'@x'$ to denote *graph isomorphism* between $canon(\gamma) \downarrow x$ and $canon(\gamma') \downarrow x'$. Both $from(\gamma)$ and $to(\gamma)$ imply $canon(\gamma) = \gamma$, so at the end of garbage collection, when the entire graph has been moved into the to-space, we will have standard isomorphism between the old graph and the new.

The main constraints satisfied by the graph are enforced in the mathematical world by the *cheney* predicate shown in Fig. 6. Additionally, the following invariant is implicit throughout the proof:

$$\texttt{to} \leqslant \texttt{scan} \leqslant \texttt{free} < \texttt{to}{+}\texttt{size} \wedge even(\texttt{from}, \texttt{to}, \texttt{scan}, \texttt{free}, \gamma)$$

Here *even* forces all objects and global pointers to be aligned on even boundaries. Notice that a graph entirely in the from-space is automatically a Cheney graph: $from(\gamma) \Rightarrow cheney(\gamma, \texttt{to}, \texttt{to})$. Similarly, if $\texttt{to} \in \gamma$ and $s = f$ then $cheney(\gamma, s, f) \Rightarrow to(\gamma)$. These observations are enough to go from the precondition to the loop invariant, and from the loop invariant to the postcondition.

Verification of `copy_ref` We omit the (simpler) spec of `copy_ref` that applies the first time it is called in `collect` (line 4) and focus instead on the calls made from the main loop (lines 8 and 11). `copy_ref` swings one field of a queued object from its original target in the from-space to its target's copy in the to-space. If the target is 0 then no action is required and the post is direct from the pre. Otherwise, we can unfold the graph (line 16) to expose the target object in the from-space. We then examine its first field, `fwd`.

If `fwd` is in the to-space, then the target object has a copy located there and we swing the pointer to it. The ramification immediately follows from Lem. 4.12, slightly modified to handle single-field updates, and updates the graph to $\gamma' = [\texttt{p} \mapsto \texttt{fwd}]\gamma$ (where a single field update $[x \mapsto y]\gamma$ corresponds to $[x \mapsto y, \gamma(x).2]\gamma$ if x is even and to $[x - 1 \mapsto \gamma(x).1, y]\gamma$ if x is odd). The actual proof effort at that point is to mathematically establish $cheney(\gamma', \texttt{p}{+}1, \texttt{free})$ and $\gamma@\texttt{to} \approx \gamma'@\texttt{to}$. The former holds because the only update is that `p` changed from queued in γ to scanned in γ' ($\texttt{p} < \texttt{p} + 1$) and from pointing to the from-space to the to-space. For the latter, notice that $canon(\gamma)(\texttt{p}) = \texttt{fwd}$, so swinging `p` to point to `fwd` gives the same canonical graph: $canon(\gamma) = canon(\gamma')$, hence $\gamma@\texttt{to} \approx \gamma'@\texttt{to}$.

If the object has not been copied yet, we reserve two units of space at the position of the `free` pointer (by advancing it, line 27), and fill them with the object's fields. Since the pool of free space is kept $*$-separated from the current graph of objects, FRAME is able to deal nicely with the heap mutations up to the assignment at line 35. Now we rewrite the state to integrate the new object into the main graph (Lem. 4.11), then swing both the current field `p` and the first field of the target object `obj` to point to the copy `new`, yielding two successive ramifications that update the global graph accordingly, which we can discharge with Lem. 4.12. Once again, RAMIFY and our library allow us to progress past updates to the shared state; the actual complexity resides in establishing mathematical facts about graphs in the postcondition. Their proof is similar to the case in which `fwd` was in the to-space to begin with. We have to prove that `new` is reachable from `to`, as required by 29, which holds because `p` is reachable from `to` and points to it. The isomorphism holds because `new` and `obj` have identical contents.

Verification of `collect` The main function first copies the root node in the graph using an alternative (simpler) spec for `copy_ref` to establish the loop invariant (line 6, in which we leave out the case $r_0 = 0$ of an empty graph). It then enters a loop that updates both fields of the first unscanned object in succession (which may queue up new objects), repeating until all objects have been scanned. The looping condition allow us to go from the invariant at line 6 to the assertion at line 7 (in particular, $\texttt{to} \rightsquigarrow^* \texttt{scan}$ by (29) so Lem. 4.10 applies). The ramification at line 9 makes interesting use of our ramification library. Lem. 4.13 tells us that each individual pointer in fromsp (as well as the other field of `scan`) is preserved. Combining this with Lem. 4.6 yields that the whole of fromsp is preserved. The graphs are updated thanks to Lem. 4.14. We finally combine both our conclusions with another application of Lem. 4.6. To deduce line 10, we fold back the sub-graph rooted at `scan` into the main one rooted at `to`, which leaves the following spatial deduction, which holds because, together, $\mathsf{graph}(\texttt{to}, \gamma')$ and fromsp contain the whole allocated heap:

$$\mathsf{graph}(\texttt{to}, \gamma') \circledast \mathsf{graph}(q_0, \gamma') \circledast \mathsf{fromsp} \vdash \mathsf{graph}(\texttt{to}, \gamma') \circledast \mathsf{fromsp}$$

The second call to `copy_ref` is analogous to the first, and after we advance `scan` we reach the loop invariant.[6]

Related work Cheney's garbage collector has been a benchmark of sorts for heap-aware verification, especially in separation logic [MAY06, TSBR08, Gas11]. Previous verifications worked by exploding the spatial graph into its individual nodes, and grouping those into several *disjoint* groups corresponding to the intersections of various heap regions (from and to-space, scanned and unscanned,

[6] In the above proof, the global variable `free` is modified by `copy_ref`, but appears in our ramified assertions. We circumvent this issue by treating `free` as a resource: we remove our knowledge about `free` when `copy_ref` is called, and only get to assume what is in `copy_ref`'s post-condition in the post-ramified state (*e.g.* line 9).

$$\dfrac{\dfrac{}{R \vdash P * (Q \mathrel{-\!\!*} R')}\ \text{Hyp.} \qquad \dfrac{\dfrac{}{\{P\}\ c\ \{Q\}}\ \text{Hyp.} \qquad \dfrac{}{\mathit{modif}(c) \cap \mathit{fv}(Q \mathrel{-\!\!*} R') = \varnothing}\ \text{Hyp.}}{\{P * (Q \mathrel{-\!\!*} R')\}\ c\ \{Q * (Q \mathrel{-\!\!*} R')\}}\ \text{Frame} \qquad \dfrac{}{Q * (Q \mathrel{-\!\!*} R') \vdash R'}\ \text{Modus Ponens}}{\{R\}\ c\ \{R'\}}\ \text{Consequence}$$

Figure 7: Proof of RAMIFY.

etc.). Our approach uses a single, *generic* inductive graph predicate, and the intricacies of reasoning about those regions is handled at the level of mathematical graphs. This division of labor yields, in our opinion, a much more pleasant and concise proof, which enjoys relatively intuitive and natural invariants.

8. Universality, Strongest Posts, and Extensions

Here we discuss the general applicability of the ramify rule as well as an alternative form of the rule. We also discuss a number of extensions to apply ramifications to more examples, including the overlapping conjunction $\circledast$, regions, and higher-order settings.

8.1 Universality of Ramification

In §3.3 we showed that the frame rule was a consequence of the ramify rule. Somewhat surprisingly, the converse is also true.

Theorem 8.1 (RAMIFY)

$$\frac{\{P\}\ c\ \{Q\} \quad R \vdash P * (Q \mathrel{-\!\!*} R')}{\{R\}\ c\ \{R'\}}\ \begin{array}{l}\mathit{fv}(Q \mathrel{-\!\!*} R') \cap \\ \mathit{modif}(c) = \varnothing\end{array}$$

Proof By the short derivation given in Fig. 7. □

Because theorem 8.1 only requires frame and consequence, ramify is valid in any separation logic. This is very handy, because it means that we do not need to modify the numerous flavors of separation logic in previous work to incorporate ramification: it has been there all along, just waiting for its importance to be recognized.

8.2 Weakest preconditions and strongest postconditions

In fact, our ramify rule appears in the separation logic folklore as a weakest precondition rule, codified as follows:

Lemma 8.1 (Weakest Pre) *Given a postcondition R' and a specification $\{P\}\ c\ \{Q\}$, then $P * (Q \mathrel{-\!\!*} R')$ is the weakest precondition,* i.e.*, given any specification $\{R\}\ c\ \{R'\}$, then $R \vdash P * (Q \mathrel{-\!\!*} R')$.*

Our examples demonstrate that we can successfully ramify with weakest precondition. Can we also succeed with strongest postcondition, *i.e.*, with the following "forward ramify" rule:

$$\text{FWRAMIFY}\quad \frac{\{P\}\ c\ \{Q\} \qquad R \vdash P * \mathsf{true} \qquad (P \mathrel{-\!\!\circledast} R) * Q \vdash R'}{\{R\}\ c\ \{R'\}}$$

The $(P \mathrel{-\!\!\circledast} R) * Q \vdash R'$ pattern is reminiscent of a pattern used in RGSep [VP07] to characterize *stability* by setting R' to R. In RGSep the focus is on concurrency, and a thread's collaborators may take an unknown number of actions. In our setting we know that a given specification will execute exactly once, which we leverage by allowing the consequent to be the more general R' rather than R. When P is precise, FWRAMIFY gives the strongest postcondition:

Lemma 8.2 (Strongest Post) *Given precondition R and $\{P\}\ c\ \{Q\}$, if P is precise then $(P \mathrel{-\!\!\circledast} R) * Q$ is the strongest postcondition.*

As it happens, whenever P is precise, RAMIFY and FWRAMIFY are each derivable from the other. However, precision is actually only needed when starting from FWRAMIFY, and so we consider RAMIFY to be fundamental. Moreover, although we were able to prove some of the examples using FWRAMIFY, we found its $\mathrel{-\!\!\circledast}$ idiom to be harder to reason about than the $\mathrel{-\!\!*}$ idiom in RAMIFY.

8.3 Extensions supporting ramification

We can ramify in any separation logic, but verifying certain programs can require various extensions, such as regions in §A. Here we detail other extensions, starting with a more careful look at $\circledast$.

The overlapping conjunction $\circledast$ Although the overlapping conjunction $\circledast$ appears occasionally in the literature (under such names as "fusion", "relevance conjunction", and "sepish"), its properties are not well-understood for abstract separation logics.

A separation algebra [COY07] is a partial commutative monoid with cancellation $(S, \oplus)$ that provides an abstract model for separation logic. Although the overlapping conjunction $\circledast$ can be defined in any separation algebra, it is not necessarily easy to use: in fact, several critical properties require stronger separation algebra axioms. We propose using a variant described by Dockins *et al.* [DHA09] that has multiple units, *disjointness* (*i.e.*, $x \oplus x = y \Rightarrow x = y$), and a kind of distributivity property called "cross split":

$$a \oplus b = z \wedge c \oplus d = z \Rightarrow \exists ac, ad, bc, bd.$$
$$ac \oplus ad = a \wedge bc \oplus bd = b \wedge ac \oplus bc = c \wedge ad \oplus bd = d$$

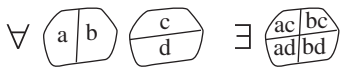

That is, if an element (*e.g.*, a heaplet) can be split in two different ways, then there are *four* subobjects which partition the original and respect the original splittings. Cross split is not discussed much in the literature, but we discovered that it is vital for reasoning about the overlapping conjunction $\circledast$, which is not even associative without it. In fact, virtually all of our proofs that use $\circledast$ assume cross split.

Many—but by no means all—separation algebras used in practice satisfy cross split, including the canonical model of heaplets as partial maps from addresses to values (quarters are found by set intersection on the domain). Users of our theory must therefore verify that the separation algebras they care about satisfy cross split.

Explicit overlapping conjunction Cherini and Blanco proposed a generalization of $P \circledast Q$ that tagged the shared core with an explicit description C [CB09], defined as follows:

$$h \vDash P \langle \circledast : C \rangle Q \stackrel{\text{def}}{=} \exists h_1, h_2, h_3.\ (h_1 \oplus h_2 \oplus h_3 = s) \wedge (h_1 \oplus h_2 \vDash P)\ \wedge\ (h_2 \vDash C)\ \wedge\ (h_2 \oplus h_3 \vDash Q)$$

This *explicit overlapping conjunction* is more expressive than $\circledast$:

$$P \langle \circledast : \mathsf{true} \rangle Q \quad \dashv\vdash \quad P \circledast Q$$

Moreover, Cherini and Blanco developed the following proof rule:

$$\text{EXPRAMIFY}\quad \frac{\{P\}\ c\ \{Q\} \qquad (C \mathrel{-\!\!\circledast} R) * C' \vdash R' \qquad Q \vdash C' * \mathsf{true}}{\{P \langle \circledast : C \rangle R\}\ c\ \{Q \langle \circledast : C' \rangle R'\}}$$

Unfortunately, EXPRAMIFY is not useful to verify any of our examples because we focus on *unspecified sharing*—that is, we do not know exactly what the overlap is (*e.g.*, the precise nodes shared between the children of a dag node), and hence cannot pick C or C' other than true. In general unspecified sharing is more difficult to verify than specified sharing, which is apparent when one tries to apply EXPRAMIFY:

$$(\mathsf{true} \mathrel{-\!\!\circledast} R) * \mathsf{true} \vdash R'$$

In other words, start from R, remove an unrestricted subheap, replace it with a second unrestricted heap, and now prove R'. Yikes!

Conversely, EXPRAMIFY cannot verify our overlaid example (§6) because instead of C and C' being too weak, they represent the entire structure (*i.e.*, $P = C = R$ and $Q = C' = R'$). Applying EXPRAMIFY then makes no progress because the "simpler" Hoare subproof $\{P\}\ c\ \{Q\}$ is actually identical to the goal.

All of that said, Cherini and Blanco demonstrate how to use EXPRAMIFY to verify programs that operate in the special case of *specified partial sharing—i.e.*, when nontrivial C and C' are known and not the entire P, Q, R, or R'. Happily, EXPRAMIFY is derivable from RAMIFY, so we can reuse all of their verifications.

Fractional shares, actions, and tight regions In §4.2 we pointed out that naïve attempts to verify `mark` using the shape-only $\mathsf{dag}(x)$ predicate were unsound. In this paper we focused on functional correctness instead, but we also experimented various other methods for guaranteeing that the graph is not overly mangled, including fractional shares [DHA09], actions in the style of RGSep [Vaf07], and a variant of regions that could prevent memory deallocation. Each method had some benefits but also required some additional formalism; the tradeoffs were unclear.

Higher-order settings In recent years there have been several flavors of separation logic to reason about higher-order state such as the resource invariants of concurrent separation logic with first-class locks [HAZ08]. Although we did not do any ramifications for genuine higher-order settings (which are often very complicated in ways unrelated to their higher-orderness), we did check a few of the ramifications from this paper in Coq within the framework of approximating separation algebras [HDA10], and believe that the higher-orderness by itself poses no fundamental difficulties.

9. Related Work

There is a large body of work, orthogonal to ours, tackling the design and proof of *algorithms* for data structures with sharing. Its counterpart in program verification spans a range of domains, and we begin this section with other separation logic based analyses.

Our reasoning about graphs owes a lot to the overlapping conjunction $\circledast$, which has roots in relevant logic [Urq72]. Many people have rediscovered it in the context of separation logic [Rey03, GMS12], who defined inductive graphs and dags as we did, but did not provide a means to reason about them. Cherini and Blanco were able to reason about a *specified* version of $\circledast$ using a more domain-specific framework than ours, as discussed in §8.3.

More recently, Mehnert *et al.* and Krishnaswami *et al.* have used some form of ramification to verify respectively implementations of snapshottable trees [MSBS12] and programs that follow the subject-observer pattern [KBA10], both of which involved unspecified sharing. Their ramifications are restricted to ad-hoc "ramification operators" tailored for each example, and the logic itself is domain-specific and done modulo a predicate on the global heap. It would be interesting to try and recast their proofs in our setting. Lee *et al.* devised an automatic analysis for threaded trees that instruments the results of separate analyses for lists and for trees [LYP11].

Moreover, several works have dealt with *definite* sharing in separation logic, *e.g.* doubly linked lists [Rey00], trees with parent pointers, skip or cyclic lists, etc. In these cases, one always knows *what* is shared and *by whom*. On the other hand, handling *indefinite* sharing, such as in this paper, was achieved only by resorting to tricks that specified or avoided the sharing. Yang's proof of the Schorr-Waite graph marking algorithm [Yan01, §7] (later mechanized in Isabelle/HOL [MN05]) does not define a spatial graph predicate, but rather describes the graph by its spanning tree. Attempts to lift this kind of reasoning to other algorithms on dags and graphs has led to convoluted predicates that explicitly deal with sharing and hack data structures into $*$-conjoined pieces, often in ways tied to the behavior of the program at hand [BCO04].

Several other frameworks have dealt with sharing in programs. In shape analysis, Hob can prove data structure consistency when one can expose a *backbone* into which objects ultimately point [WKL+06], and TVLA has been used to prove partial correctness of a mark-and-sweep garbage collector and the Schorr-Waite algorithm [MSRF04]. Hawblitzel and Petrank have used Boogie to automatically verify garbage collectors [HP09]. However, these works do not provide *compositional* reasoning for sharing.

It would be interesting to see if we can import ramification into other frameworks, such as Dafny [Lei10], whose reasoning about the heap is based on dynamic frames (a cousin of separation logic).

10. Conclusion

We have presented a new paradigm, ramification, valid in any separation logic, for the compositional verification of programs that manipulate data structures with both specified and unspecified sharing. We gave a ramification library that helps simplify ramification entailments in general and reduces local spatial updates to abstract mathematical reasoning. We have demonstrated the applicability of our framework by providing concise, local specifications for a range of examples and data structures, including Cheney's garbage collector. These initial successes lead us to believe that ramification provides a robust basis for elegant, compositional reasoning about sharing in data structures.

Acknowledgments

We deeply thank Peter O'Hearn for his continuous help and encouragement. We also benefited from discussions with Josh Berdine, Richard Bornat, Gareth Smitch, David Walker, Hongseok Yang, and especially with Matthew Parkinson, who first suggested that our initial semantic account of ramification was expressible as a separation logic entailment. Finally, we thank the anonymous reviewers for their suggestions and enthusiasm.

This research was supported by a Lee Kuan Yew Postdoctoral Fellowship and EPSRC Programme Grant "Resource Reasoning".

References

[BCO04] R. Bornat, C. Calcagno, and P. O'Hearn. Local reasoning, separation and aliasing. In *SPACE*, 2004.

[BCY06] R. Bornat, C. Calcagno, and H. Yang. Variables as resource in separation logic. *ENTCS*, 155, 2006.

[Bor00] R. Bornat. Proving pointer programs in Hoare logic. In *MPC*, 2000.

[CB09] R. Cherini and J. O. Blanco. Local reasoning for abstraction and sharing. In *SAC*, 2009.

[Che70] C. J. Cheney. A nonrecursive list compacting algorithm. *C. ACM*, 13(11), 1970.

[COY07] C. Calcagno, P. W. O'Hearn, and H. Yang. Local action and abstract separation logic. In *LICS*, 2007.

[DHA09] R. Dockins, A. Hobor, and A. W. Appel. A fresh look at separation algebras and share accounting. In *APLAS*, 2009.

[Fin87] J. Finger. *Exploiting constraints in design synthesis*. PhD thesis, Stanford University, 1987.

[Gas11] H. Gast. Developer-oriented correctness proofs - a case study of Cheney's algorithm. In *ICFEM*, 2011.

[GMS12] P. Gardner, S. Maffeis, and G. D. Smith. Towards a program logic for JavaScript. In *POPL*, 2012.

[HAZ08] A. Hobor, A. W. Appel, and F. Zappa Nardelli. Oracle semantics for concurrent separation logic. In *ESOP*, 2008.

[HDA10] A. Hobor, R. Dockins, and A. W. Appel. A logical mix of approximation and separation. In *APLAS*, ENTCS, 2010.

[HP09] C. Hawblitzel and E. Petrank. Automated verification of practical garbage collectors. In *POPL*, 2009.

[IO01] S. S. Ishtiaq and P. W. O'Hearn. BI as an assertion language for mutable data structures. In *POPL*, 2001.

[KBA10] N. Krishnaswami, L. Birkedal, and J. Aldrich. Verifying event-driven programs using ramified frame properties. In *TLDI*, 2010.

[Lei10] K. R. M. Leino. Dafny: An automatic program verifier for functional correctness. In *LPAR*, 2010.

[LG88] J. M. Lucassen and D. K. Gifford. Polymorphic effect systems. In *POPL*, 1988.

[LYP11] O. Lee, H. Yang, and R. Petersen. Program analysis for overlaid data structures. In *CAV*, 2011.

[MAY06] N. Marti, R. Affeldt, and A. Yonezawa. Formal verification of the heap manager of an operating system using separation logic. In *ICFEM*, 2006.

[MN05] F. Mehta and T. Nipkow. Proving pointer programs in higher-order logic. *Inf. Comput.*, 199(1-2), 2005.

[MSBS12] H. Mehnert, F. Sieczkowski, L. Birkedal, and P. Sestoft. Formalized verification of snapshotable trees: Separation and sharing. In *VSTTE*, 2012.

[MSRF04] R. Manevich, S. Sagiv, G. Ramalingam, and J. Field. Partially disjunctive heap abstraction. In *SAS*, 2004.

[Rey00] J. C. Reynolds. Intuitionistic reasoning about shared mutable data structure. In *Millennial Perspectives in Computer Science*, Cornerstones of Computing, 2000.

[Rey02] J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *LICS*, 2002.

[Rey03] J. C. Reynolds. A short course on separation logic. `http://www.cs.cmu.edu/afs/cs.cmu.edu/project/fox-19/member/jcr/wwwaac2003/notes7.ps`, 2003.

[Thi01] M. Thielscher. The qualification problem: A solution to the problem of anomalous models. *Artificial Intelligence*, 131(1), 2001.

[TSBR08] N. Torp-Smith, L. Birkedal, and J. C. Reynolds. Local reasoning about a copying garbage collector. *ACM TOPLAS*, 30(4), 2008.

[Urq72] A. Urquhart. Semantics for relevant logics. *J. Symb. Log.*, 37(1), 1972.

[Vaf07] V. Vafeiadis. *Modular fine-grained concurrency verification*. PhD thesis, University of Cambridge, 2007.

[VP07] V. Vafeiadis and M. J. Parkinson. A marriage of rely/guarantee and separation logic. In *CONCUR*, 2007.

[WKL+06] T. Wies, V. Kuncak, P. Lam, A. Podelski, and M. C. Rinard. Field constraint analysis. In *VMCAI*, 2006.

[Yan01] H. Yang. *Local Reasoning for Stateful Programs*. PhD thesis, University of Illinois, 2001.

A. Copying Dags

The program in Fig. 8 makes a deep (structure-preserving) copy of a dag, using a data field in each original node to record the location of its copy when it exists (or 0 otherwise). Initially, all the copy fields of $\mathsf{dag}(\mathtt{x})$ must be set to 0, and at the end all the nodes reachable from `x` will have been copied into a new dag whose root is returned by `copy_dag`. In the intermediate recursive calls, parts of the dag rooted at the argument will have already been copied.

To cut down on the amount of formalism we must present to verify a reasonable-looking specification for `copy_dag`, we will utilize regions in the following ad-hoc way. We assume two regions, α and β, bound at the "top level", in the meta context. The initial dag must come in region α, and `malloc` always allocates in region β; afterwards the new copy will be in β and the original will remain in α. This specification is not as general as one would want, since *e.g.* it prevents us from verifying the copying of a copy; a far better specification would take the regions as parameters, but all of the extra wizardry would be in the (orthogonal) region management system rather than in the ramifications. Ideally, such a system would have features such as general-purpose region creation, destruction, and merging, as well as a good handle on region variable scoping.

Describing completed and in-process dag copies We represent an entirely copied dag $\delta = (V, D, L, E)$ rooted at x and its copy rooted at y by the predicate $\mathsf{ddag}(x, y, \delta)$ (or *double dag*):

$$\mathsf{ddag}(x,y,\delta) \stackrel{\text{def}}{=} (x = y = 0 \wedge \mathsf{emp}) \vee (\mathsf{dag}_\alpha(x,\delta) * \mathsf{dag}_\beta(y, copy(\delta)) \wedge L(x) = y \wedge y \neq 0)$$

The nodes in the first dag are described by the graph δ. Because we store the addresses of the copy in the data fields, δ is also enough to describe the copy via $copy(\delta) = (V', D, L', E')$, where

$$V' = \{v' \mid \exists v \in V.\, L(v) = v' \wedge v' \neq 0\}$$

$$L'(v) = 0$$

$$E'(v) = (l', r') \text{ if } \begin{cases} \exists v' \in V.\, \delta(v') = (v, l, r) \wedge \\ (l = l' = 0 \vee L(l) = l') \wedge \\ (r = r' = 0 \vee L(r) = r') \end{cases}$$

The predicate ddag describes the postcondition for `copy_dag`; our next task is to define the precondition. Because some parts of the dag may have already been copied, the *in-copy dag* predicate $\mathsf{icdag}(x, \delta)$ describes a single dag in region α and a *set* of dags in region β corresponding to any previously copied sub-dags.

$$\mathsf{icdag}(x,\delta) \stackrel{\text{def}}{=} \mathsf{dag}_\alpha(x,\delta) * \mathsf{dags}_\beta(pr(x,\delta), copy(\delta))$$

$pr(x, \delta)$ (processed roots) finds the roots of the copied sub-dags:

$$pr(x,\delta) \stackrel{\text{def}}{=} \begin{cases} \varnothing & \text{if } x = 0 \\ pr(l,\delta) \cup pr(r,\delta) & \text{if } \delta(x) = (0,l,r) \\ \{x\} & \text{otherwise} \end{cases} \quad (32)$$

Observe that when x is copied, *i.e.* $\delta(x) = (y, l, r)$ and $y \neq 0$, then

$$\mathsf{icdags}(x,\delta) \dashv\vdash \mathsf{ddags}(x,y,\delta) \quad (33)$$

We will use this equivalence to move between the precondition and the postcondition when we discover that the dag is already copied.

When we wish to reason entirely about the copies we write $\mathsf{cdags}(x, \delta)$ (*i.e., copy dags*) for $\mathsf{dags}(pr(x, \delta), copy(\delta))$. Note that if x is not yet copied, *i.e.* $\delta(x) = (0, l, r)$, then, using the second case in the definition of pr, we deduce that

$$\mathsf{cdags}(x,\delta) \dashv\vdash \mathsf{cdags}(l,\delta) \circledast \mathsf{cdags}(r,\delta), \text{ and thus}$$

$$\mathsf{icdags}(x,\delta) \dashv\vdash x \mapsto_\alpha 0, l, r * (\mathsf{icdags}(l,\delta) \circledast \mathsf{icdags}(r,\delta)) \quad (34)$$

Finally, to reflect the fact that already copied parts of the dag will not be changed by `copy_dag`, we define the relation $\delta' \geqslant_x \delta$ between two dags $\delta = (V, D, L, E)$ and $\delta' = (V, D, L', E)$, true when $\delta' \uparrow x = \delta \uparrow x$ and $\delta' \downarrow x$ is "more copied" than $\delta \downarrow x$:

$$\forall v \in reach(\delta, x).\quad \begin{array}{l}(L(v) \neq 0 \Rightarrow L'(v) = L(v)) \wedge \\ (L'(v) = 0 \Rightarrow L(v) = 0)\end{array}$$

We will write $\delta' \geqslant \delta$ when $\exists x.\, \delta' \geqslant_x \delta$.

```
1  struct node {struct node *c,*l,*r;};
2  struct node *
3  copy_dag(struct node *x) { // {icdag(x, δ)}
4    struct node *l,*r,*ll,*rr,*y;
5    if (!x) return 0;
6    if (x->c) return x->c;
7    l = x->l; r = x->r;
8    y = malloc(sizeof(struct node));
9    x->c = y;
10 // { x ↦α y,l,r * (icdag(l, δ) ⊛ icdag(r, δ)) * y ↦β −,−,− ∧ δ(x) = (0,l,r) }
11   ll = copy_dag(l);
12 //↯(35) { x ↦α y,l,r * (ddag(l,ll,δ′) ⊛ icdag(r,δ′)) * y ↦β −,−,− ∧ δ(x) = (0,l,r) ∧ δ′ ⩾1 δ }
13   rr = copy_dag(r);
14 //↯(36) { x ↦α 0,l,r * (ddag(l,ll,δ″) ⊛ ddag(r,rr,δ″)) * y ↦β −,−,− ∧ δ(x) = (0,l,r) ∧ δ′ ⩾1 δ ∧ δ″ ⩾r δ′ }
15   y->c = 0; y->l = ll; y->r = rr;
16 // { x ↦α y,l,r * y ↦β 0,ll,rr * (ddag(l,ll,δ″) ⊛ ddag(r,rr,δ″)) ∧ δ(x) = (0,l,r) ∧ δ′ ⩾1 δ ∧ δ″ ⩾r δ′ }
17   return y;
18 } // {ddag(x, y, δ‴) ∧ δ‴ ⩾x δ}
```

Figure 8: Proof sketch of dag copy.

Verification of `copy_dag` Now we annotate the program in Fig. 8 with the key assertions to prove the following specification:

$$\{\mathsf{icdag}(x,\delta)\}\ \texttt{y = copy\_dag(x)}\ \{\mathsf{ddag}(x,y,\delta') \wedge \delta' \geqslant_x \delta\}$$

If the dag is empty then the postcondition is trivially satisfied (line 5); if the node has already been copied (line 6) then equation 33 yields the postcondition. The real meat of the algorithm is in the ramifications from the two recursive call sites and the entailment of the postcondition from line 16. The two ramifications are as follows:

$$\begin{array}{r}\mathsf{icdag}(\mathtt{l},\delta) \circledast \mathsf{icdag}(\mathtt{r},\delta)\\ \vdash \mathsf{icdag}(\mathtt{l},\delta) * (\mathsf{ddag}(\mathtt{l},\mathtt{ll},\delta') \wedge \delta' \geqslant_{\mathtt{l}} \delta \mathrel{-\!\!*}\\ \mathsf{ddag}(\mathtt{l},\mathtt{ll},\delta') \circledast \mathsf{icdag}(\mathtt{r},\delta') \wedge \delta' \geqslant_{\mathtt{l}} \delta)\end{array} \tag{35}$$

$$\begin{array}{r}\mathsf{ddag}(\mathtt{l},\delta') \circledast \mathsf{icdag}(\mathtt{r},\delta')\\ \vdash \mathsf{icdag}(\mathtt{r},\delta') * (\mathsf{ddag}(\mathtt{r},\mathtt{rr},\delta'') \wedge \delta'' \geqslant_{\mathtt{r}} \delta' \mathrel{-\!\!*}\\ \mathsf{ddag}(\mathtt{l},\mathtt{ll},\delta'') \circledast \mathsf{ddag}(\mathtt{r},\mathtt{rr},\delta'') \wedge \delta'' \geqslant_{\mathtt{r}} \delta')\end{array} \tag{36}$$

As with `mark`, the second ramification follows from the first by swapping the roles of `r` and `l` and observing that when $\delta' \geqslant \delta$

$$\mathsf{ddag}(x,y,\delta) \vdash \mathsf{icdag}(x,\delta') \Leftrightarrow \mathsf{ddag}(x,y,\delta') \Leftrightarrow \mathsf{ddag}(x,y,\delta)$$

Regions let us split the first ramification (35) using the Lem. 4.5 from our ramification library, yielding two simpler ramifications in which $\delta' \geqslant_{\mathtt{l}} \delta$, and, by the definition of ddag, $\mathtt{l} = \mathtt{ll} = 0 \vee \delta'(\mathtt{l}) = (\mathtt{ll},-,-)$. The first half of (35), in region α,

$$\begin{array}{r}\mathsf{dag}(\mathtt{l},\delta) \circledast \mathsf{dag}(\mathtt{r},\delta) \vdash \mathsf{dag}(\mathtt{l},\delta) * (\mathsf{dag}(\mathtt{l},\delta') \mathrel{-\!\!*}\\ \mathsf{dag}(\mathtt{l},\delta') \circledast \mathsf{dag}(\mathtt{r},\delta')),\end{array} \tag{37}$$

is direct from Lem. 4.14. The second half of (35), in region β, is

$$\begin{array}{r}\mathsf{cdags}(\mathtt{l},\delta) \circledast \mathsf{cdags}(\mathtt{r},\delta) \vdash \mathsf{cdags}(\mathtt{l},\delta) * (\mathsf{dag}(\mathtt{ll},\delta'_c) \mathrel{-\!\!*}\\ \mathsf{dag}(\mathtt{ll},\delta'_c) \circledast \mathsf{cdags}(\mathtt{r},\delta'))\end{array} \tag{38}$$

where $\delta'_c = copy(\delta')$. This ramification is more involved because the copied roots of δ' starting from `r` may differ from the previous ones in δ. Instantiating Lem. 4.14 with $S_1 = pr(\delta,\mathtt{l})$, $S_2 = pr(\delta,\mathtt{r})$ and $S'_1 = \{\mathtt{ll}\}$ yields this entailment, which is only halfway there, because it features the sub-dags rooted at $pr(\delta,\mathtt{r})$, whereas we want those rooted at $pr(\delta',\mathtt{r})$:

$$\begin{array}{r}\mathsf{cdags}(\mathtt{l},\delta) \circledast \mathsf{cdags}(\mathtt{r},\delta) \vdash \mathsf{cdags}(\mathtt{l},\delta) *\\ (\mathsf{dag}(\mathtt{ll},\delta'_c) \mathrel{-\!\!*} \mathsf{dag}(\mathtt{ll},\delta'_c) \circledast \mathsf{dags}(pr(\delta,\mathtt{r}),\delta'_c))\end{array}$$

To complete this proof, we remark that the copied roots of `r` in δ' and in δ satisfy the following relations, hence Lem. 4.10 applies:

$$\begin{array}{ll} pr(\delta,\mathtt{r}) \subseteq reach(\delta'_c, pr(\delta',\mathtt{r})) & (\delta' \geqslant \delta)\\ pr(\delta',\mathtt{r}) \subseteq pr(\delta,\mathtt{r}) \cup reach(\delta'_c,\mathtt{ll}) & (\delta' \geqslant_{\mathtt{l}} \delta)\end{array}$$

To reach the postcondition from line 16, the sub-copies on each side of the overlapping conjunction need to be disentangled from the original sub-dags using regions and equation 25 in the following derivations, where $\delta(\mathtt{x}) = (0,\mathtt{l},\mathtt{r})$, $\delta' \geqslant_{\mathtt{l}} \delta$, and $\delta'' \geqslant_{\mathtt{r}} \delta'$:

$$\begin{array}{l}\mathtt{x} \mapsto_\alpha \mathtt{y},\mathtt{l},\mathtt{r} * \mathtt{y} \mapsto_\beta 0,\mathtt{ll},\mathtt{rr} *\\ (\mathsf{dag}_\alpha(\mathtt{l},\delta'') * \mathsf{dag}_\beta(\mathtt{ll}, copy(\delta''))) \circledast\\ (\mathsf{dag}_\alpha(\mathtt{r},\delta'') * \mathsf{dag}_\beta(\mathtt{rr}, copy(\delta')))\\ \vdash \mathsf{dag}_\alpha(\mathtt{x},\delta''') * \mathsf{dag}_\beta(\mathtt{y}, copy(\delta''')) \wedge \delta''' = [\mathtt{x} \mapsto (\mathtt{y},\mathtt{l},\mathtt{r})]\delta''\\ \vdash \mathsf{ddag}(\mathtt{x},\mathtt{y},\delta''') \wedge \delta''' \geqslant_{\mathtt{x}} \delta\end{array} \tag{39}$$

The last deduction step uses this mathematical fact:

$$\delta(x) = (0,l,r) \wedge \delta'(x) = (y,l,r) \wedge \delta' \geqslant_l \delta \wedge \delta' \geqslant_r \delta \Rightarrow \delta' \geqslant_x \delta$$

B. Disposing a Graph

Let us show how to verify the depth-first search spanning tree procedure for binary graphs, as presented in Fig. 9.

The desired top-level specification for `spanning` is that, starting from an unmarked graph γ, we remove some edges (indicated by the predicate $\sqsubseteq$) and get a tree τ that covers the same set of nodes:

$$\begin{array}{l}\{\mathsf{graph}(x,\gamma) \wedge unmarked(\gamma)\}\\ \texttt{spanning(x)}\\ \{\mathsf{tree}(x,\tau) \wedge \tau \sqsubseteq \gamma \wedge reach(\tau,x) = reach(\gamma,x)\}\end{array}$$

The predicate $(V,D,L,E) \sqsubseteq (V',D',L',E')$ is true when $(V',D',L') = (V,D,L)$, and E has "fewer edges" than E':

$$\forall v \in V.\, E'(v) = (l',r') \Rightarrow E(v) = (l,r) \wedge l \in \{l',0\} \wedge r \in \{r',0\}$$

```
1  void spanning(struct node *x) {
2  // {graph(x,γ) ∧ γ(x) = (0,−,−)}
3    struct node *l,*r;
4    l = x->l; r = x->r;
5  // {x ↦ 0,l,r ⊛ graph(l,γ) ⊛ graph(r,γ) ∧ γ(x) = (0,l,r)}
6    x->m = 1;
7  //↯ { x ↦ 1,l,r ⊛ graph(l,γ1) ⊛ graph(r,γ1) ∧
       γ(x) = (0,l,r) ∧ γ1 = m1(γ,x) }
8    if (l && !l->m)
9  // { x ↦ 1,l,r ⊛ graph(l,γ1) ⊛ graph(r,γ1) ∧
       γ(x) = (0,l,r) ∧ γ1 = m1(γ,x) ∧ γ1(l) = (0,−,−) }
10     spanning(l);
11   else x->l = 0;
12 //↯ { (x ↦ 1,l,r ⊛ tree(l,γ2) ⊛ graphs(pr(l,γ1),γ2) ⊛
       graph(r,γ2) ∧ γ(x) = (0,l,r) ∧ γ1 = m1(γ,x) ∧
       γ1(l) = (0,−,−) ∧
       γ2 ⊑ m(γ1,l) ∧ reach(γ2,l) = reach0(γ1,l)) ∨ ··· }
13   if (r && !r->m) spanning(r);
14   else x->r = 0;
15 //↯ { ( x ↦ 1,l,r ⊛ tree(l,γ3) ⊛ graphs(pr(l,γ1),γ3) ⊛
       tree(r,γ3) ⊛ graphs(pr(r,γ2),γ3) ∧ γ(x) = (0,l,r) ∧
       γ1 = m1(γ,x) ∧ γ1(l) = (0,−,−) ∧ γ2(r) = (0,−,−) ∧
       γ2 ⊑ m(γ1,l) ∧ reach(γ2,l) = reach0(γ1,l) ∧
       γ3 ⊑ m(γ2,r) ∧ reach(γ3,r) = reach0(γ2,r) )
     ∨ ( x ↦ 1,l,0 ⊛ tree(l,γ'3) ⊛ graphs(pr(l,γ1),γ'3) ⊛
       graph(r,γ'3) ∧ γ(x) = (0,l,r) ∧ γ1 = m1(γ,x) ∧
       γ1(l) = (0,−,−) ∧ (γ2(r) = (1,−,−) ∨ r = 0) ∧
       γ2 ⊑ m(γ1,l) ∧ reach(γ2,l) = reach0(γ1,l) ∧
       γ'3 = [x ↦ (1,l,0)]γ2 )
     ∨ ··· }
16 } // { tree(x,γ') ⊛ graphs(pr(x,γ),γ') ∧
       γ' ⊑ m(γ,x) ∧ reach0(γ,x) = reach(γ',x) }
```

Figure 9: Spanning tree of a binary graph.

During the execution, the graph will be partially marked, and the effect of `spanning` on such graphs is thus subtler. Assuming that the root x of the graph γ has not been marked yet, it transforms the *unmarked* part of γ that is reachable from x (it was not apparent in the top-level specification, wherein $reach_0(\gamma,x) = reach(\gamma,x)$) into a tree covering the same nodes, overlapped with some subgraphs. As seen in Fig. 9, these extra subgraphs are precisely those that start at a marked node reachable from x via unmarked nodes in the original graph, using the pr predicate (32) from Apx. A.

The proof of `spanning` has four main branches, corresponding to whether each of the left and right sub-graphs has to be examined or not (notice that `spanning` assumes a non-empty graph as a precondition). In Fig. 9, we only show the proof sketch corresponding to the case where the left sub-graph was non-empty and unmarked. Marking the root `x` is done as in `mark` for graphs. To handle the recursive call of line 10, we have to prove the following ramification:

$$\begin{array}{r}\mathsf{graph}(\mathtt{l},\gamma_1) \circledast \mathsf{graph}(\mathtt{x},\gamma_1) \vdash \mathsf{graph}(\mathtt{l},\gamma_1) *\\ (\mathsf{tree}(\mathtt{l},\gamma_2) \circledast \mathsf{graphs}(pr(\mathtt{l},\gamma_1),\gamma_2) \mathrel{-\!\!*}\\ \mathsf{tree}(\mathtt{l},\gamma_2) \circledast \mathsf{graphs}(pr(\mathtt{l},\gamma_1),\gamma_2) \circledast \mathsf{graph}(\mathtt{x},\gamma_2))\end{array}$$

Rewriting $\mathsf{tree}(\mathtt{l},\gamma_2)$ as $\mathsf{graph}(\mathtt{l},\gamma_2) \wedge (\gamma_2 \downarrow \mathtt{l}$ is a tree), using an analogue of Lem. 4.8 for trees turns this ramification into an application of Lem. 4.14. The same trick can be used to obtain the first disjunct of line 15 (corresponding to the case where `spanning(r)` was applied), while the second disjunct is an application of Lem. 4.12.

Because the nodes covered by the sub-trees at `l` and `r` are marked and form the same set as the nodes reachable via unmarked paths in the graphs before each recursive call, we can disentangle both trees and the roots in the first disjunct of line 15 to form the spatial part

$$\begin{array}{c}(\mathtt{x} \mapsto 1,\mathtt{l},\mathtt{r} * \mathsf{tree}(\mathtt{l},\gamma_3) * \mathsf{tree}(\mathtt{r},\gamma_3))\\ \circledast \mathsf{graphs}(pr(\mathtt{l},\gamma_1),\gamma_3) \circledast \mathsf{graphs}(pr(\mathtt{r},\gamma_2),\gamma_3)\end{array}$$

The post follows from the pure facts. The other disjuncts are similar.