

Local reasoning, separation and aliasing

Richard Bornat
School of Computing Science
Middlesex University.
R.Bornat@mdx.ac.uk

Cristiano Calcagno
Department of Computing
Imperial College
University of London
ccris@doc.ic.ac.uk

Peter O'Hearn
Department of Computer Science
Queen Mary College
University of London
ohearn@dcs.qmul.ac.uk

Abstract

Structures built by pointer aliasing, such as DAGs and graphs, are notoriously tricky to deal with. The mechanisms of separation logic can deal with these structures, but so far this has been done by the maintenance of a global invariant. Specifications and proofs which use local reasoning, and which may point the way to a structured programming for pointers, are discussed. An idiom for inclusion sharing, where one structure is included in another, is presented. A notion of ‘partial graphs’ – graphs with dangling pointers – is used to facilitate proof.

1 Introduction

The purpose of a formal logic is to codify reasoning, to enable simple mechanical steps that correspond to reliable but intricate reasoning in a semantic domain. It's reasonable to demonstrate that a novel logic helps us to solve old problems in new ways. A sound logic, however, will prove nothing that is not already a semantic consequence so in a sense it adds nothing to knowledge: all that you can prove with it you could already prove without it. In that sense it's futile to demonstrate its utility, but it's still important to try, because ease of reasoning matters.

The notion of structured programming is inextricably intertwined with the properties of formal logics for reasoning about programs. Proofs of structured programs come apart at the same seams as the programs themselves, and compose likewise. Build your programs from these components in this way, the message goes, and it will be easier to build a correct program. It is so effective that nowadays it's hard to find programmers who don't write structured programs. Designers of novel logics can quite reasonably be asked if their work will lead to similarly powerful reasoning and similarly effective advice.

Separation logic [9, 4, 5, 8], which derives from Bunched Implication (BI) logic [6, 7], aims to simplify reasoning about programs that mutate data structures in a heap – pointer programs, for short.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Pointers have been thought of as the gotos of data structuring, and so they can be, but if we can organise things so that our data structures and our programs come apart neatly in similar ways – structured reasoning for the heap – we shall perhaps be able to use pointers reliably. Separation logic facilitates *local reasoning* about parts of data structures: if an assignment or a procedure call alters only one corner of one data structure, then we can easily express and exploit the fact that it doesn't alter any other part of it nor any part of any separate data structure. The advantages can be seen most clearly in Yang's proof of the Schorr-Waite algorithm [12, 13], which elegantly avoids the struggles of [1] to show the same result using Hoare logic and array assignment. Apart from that work and Torp-Smith's proof of the Cheney algorithm [11], most work using separation logic has focussed on lists and trees, where separation is easily described.

In this paper we discuss three formal proofs (alas, not yet mechanically checked) of algorithms which deal with internal sharing of structure; two of the proofs deal with aliasing as well. Our aim has been to make the reasoning as local as possible. Although we can see considerable possibilities for further improvement, we can already claim significant success in extending the range of separation logic to new and difficult areas.

2 A brief introduction to separation logic

Assignment to a variable affects just that variable. Hoare's variable assignment rule

$$\{R_E^x\} x := E \{R\}$$

elegantly captures that property in a substitution that affects only one variable. The rule supports local reasoning because substitution has no effect on any other variable. Assignment to an array element affects just that element, but Hoare's array element assignment rule

$$\{R_{a \oplus (l:E)}^a\} a[l] := E \{R\}$$

isn't correspondingly local. Instead, the assignment affects every formula that mentions the array, not merely those which mention the assigned element. As a consequence, the difficulty of reasoning in Hoare logic about programs that use arrays revolves around the arithmetic of array indexing. A program heap is a giant array, and is usually sparse. Reasoning about heaps in Hoare logic is just too hard; if we are to do the job, we must find an easier way.

A program's heap contains cells addressed by integers (not necessarily a contiguous range of addresses, nor necessarily starting at 0). Heap cells contain values. In particular, they may contain integers which should be read as pointers, and thus we can have recursive and cyclic data structures.

Heap space can be acquired and, once acquired, released. The heap a program is allowed to use is limited to what was provided when execution starts plus what was acquired during execution and not yet released. It follows that only some integers point into the heap, while others are ‘dangling pointers’. Dereferencing a dangling pointer is always a mistake and can be a disaster; for simplicity we suppose that it always crashes the program. Avoiding such crashes is one of the practical difficulties of programming with pointers. Dangling pointers are not a difficulty for separation logic, however: there’s a real sense in which they are the point, the engine that makes the reasoning tick.

A program has a ‘stack’ of named variables (think: collection of registers) outside the heap. There are no address operations, implicit or explicit, on the stack, so aliasing within it is impossible. Variables may contain integers and thus point into the heap (or dangle).

Our programming language is ‘while’ programs with recursive procedures, using a restricted range of assignment instructions, intended to correspond closely to the operation of a conventional von Neumann machine. We use Hoare triples $\{Q\}C\{R\}$ to specify partial correctness: program C , started in a state (heap + stack) which satisfies Q , is guaranteed (i) not to crash by referencing a heap cell outside its footprint and (ii) if it terminates, to arrive at a state satisfying R . The novelty of the separation logic approach is the interpretation of formulas describing the heap in Q and R .

The simplest heap is **emp**, which contains no cells and allows no dereferencing at all. A single-cell heap, for example $x \mapsto V$, allows dereferencing and update of just one location. In this heap $[x]$ (read ‘contents of x ’) gives the value in the heap cell indexed by x , and delivers V ; the instruction $[x] := E$ (see below) updates that same cell, and after it the heap is $x \mapsto E$. The construction $x \mapsto _$ is a useful shorthand for $\exists v. x \mapsto v$. The intuitionistic operator \hookrightarrow doesn’t define a heap but allows searching: $x \hookrightarrow E$ means the current heap has a sub-heap which satisfies $x \mapsto E$. In examples we write $x \not\hookrightarrow _$ as shorthand for $\neg(\exists v. x \hookrightarrow v)$ – i.e. x doesn’t point anywhere in the current heap.

The (\star) connective deals with *disjoint* heaps (heaps with disjoint address domains). $A \star B$ holds of a heap just when it can be divided into two disjoint heaps, where simultaneously A holds in one and B in the other. (\star) is associative and commutative; in BI it’s a multiplicative version of conjunction. As you would expect, $A \star \mathbf{emp} = \mathbf{emp} \star A = A$. You can deduce disjointness of address domains when heaps are combined: for example $(x \mapsto _ \star y \mapsto _) \vdash x \neq y$, and $x \mapsto V \star x \mapsto W$ must be false whatever the values of V and W , since the integer x can’t point to two separate cells at the same time. Note that disjoint heaps may contain dangling pointers across the divide, so $x \mapsto y \star y \mapsto x$ is acceptable.

Address arithmetic is allowed, since pointers are really just integers. In particular the construction $x \mapsto V_1, V_2, \dots, V_n$ is shorthand for $x \mapsto V_1 \star x + 1 \mapsto V_2 \star \dots \star x + n - 1 \mapsto V_n$. In our programs we are content to use field-access formulae $x.\mathit{field}$ as shorthand for $[x + K]$ for some agreed fixed integer K .

The (\wedge) connective deals with *identical* heaps. $A \wedge B$, normal additive conjunction, holds of a heap if A holds of the whole, and similarly B . In proofs below we make use of the construction $(A \star \mathbf{True}) \wedge B$: \mathbf{True} holds of any heap, so this says that A holds in part of the heap, and B holds in the whole, making a kind of subheap relation. We can also use $A \wedge B$ if either A or B is ‘pure’, saying nothing about the heap at all. In examples we give (\wedge) a lower syntactic priority than (\star) , so that $A \star B \wedge C \wedge D$ is read as

$$(A \star B) \wedge C \wedge D.$$

(\multimap) is the ‘magic wand’ connective. $A \multimap B$ is true of a heap if, whenever it is joined to a disjoint heap which makes A true, B holds in the result. In the particular case that there is only one heap which satisfies A , $A \multimap B$ describes a B -heap with an A -shaped hole. It is a multiplicative form of implication, and $A \star (A \multimap B) \vdash B$.

In $\{Q\}C\{R\}$, C starts in a state satisfying Q and generates a state satisfying R . Q describes the whole heap used by C , unless it gets some more from a resource-generator like ‘new’. There are also resource-destructors like ‘dispose’, and thus it follows that R may describe a larger or a smaller heap than Q does. Because of this resource interpretation, and in contrast to our conventional treatment of assignment to stack variables, we can read Q as circumscribing C ’s initial footprint in the heap.

The most important inference rule is the Frame Rule, which allows us to localise reasoning about resources (it requires a proviso that C doesn’t assign stack variables free in P because of our conventional treatment of stack-variable assignment).

$$\frac{\{Q\}C\{R\}}{\{P \star Q\}C\{P \star R\}} \quad (\text{modifies } C \cap \text{vars } P = \emptyset)$$

If C doesn’t go outside the heap described by Q , and if P describes a disjoint heap, then clearly C leaves the heap described by P unchanged. In effect we can pull Q out of the precondition, convert it to R with C , and plug R back in to the postcondition.

We recognise three forms of assignment. It may seem irksome to have a limited collection, but it’s easy to see that more complicated constructions can be compiled down into the permitted forms by using extra assignments to stack variables. Assignment to stack variables is treated as in ordinary Hoare logic. In these rules E is a ‘pure’ expression which doesn’t refer to the heap. Because of the frame rule, assignments that affect the heap can be defined by giving a minimum resource context. (The naming restrictions on variable assignment can be overcome with \exists bindings of initial values.)

$$\begin{array}{l} \{R_E^x\} \quad x := E \quad \{R\} \\ \{x \mapsto _ \} [x] := E \quad \{x \mapsto E\} \\ \{E' \mapsto E\} \quad x := [E'] \quad \{E' \mapsto E \wedge x = E'\} \quad (x \text{ not free in } E, E') \end{array}$$

We can treat directly the kind of primitive storage allocation which underlies more sophisticated forms.

$$\begin{array}{l} \{\mathbf{emp}\} \quad x := \mathbf{new}(E) \quad \{x \mapsto E\} \\ \{E \mapsto _ \} \quad \mathbf{dispose } E \quad \{\mathbf{emp}\} \end{array}$$

These pre- and post-conditions give minimum resource contexts. Together with the frame rule they tell us that new delivers a pointer which is distinct from any in use in any other context (but not necessarily distinct from any other which may have been disposed in the past, nor from any integer values which may be lying around in stack or heap) and that dispose denies us future access to a currently-accessible cell. In algorithms we write $\mathbf{dispose } E \ N$ as a shorthand for $\mathbf{dispose } E; \mathbf{dispose}(E + 1); \dots; \mathbf{dispose}(E + N - 1)$.

In reasoning about procedures we follow [5]. Our procedures don’t alter any of their arguments, and they have no free variables, though of course they may have side-effects on the heap. As a notational convenience our procedures can return a result which can be assigned, but in every case our code could be transcribed into a more conventional version which uses a result parameter. Where we have

written an assignment to a variable that isn't in the procedure's parameter list – $l := \dots$ in (9), for example – the cautious reader should insert a local variable definition. In dealing with recursive steps, we have appealed to an instance of the specification as an inductive hypothesis. We haven't given a proof of termination of any of the examples, but structural induction arguments based on the formula that is represented in the heap (see section 4 below) would seem to be straightforward.

3 Separation and aliasing

The major practical difficulty of reasoning in a conventional Hoare logic about programs that mutate arrays is array-index aliasing. Different index formulae denote the same array element, and formulae change their reference as the values of variables change. Program proofs must therefore focus on index formula arithmetic, which is difficult to deal with logically.

The message of separation logic is that in dealing with programs that mutate data structures interconnected by pointers, separation trumps aliasing. The list, tree and graph algorithms investigated in [1], for example, depend on separation between structures. To verify those algorithms with array-index reasoning, it's necessary to describe that separation in predicate logic and to laboriously prove after each algorithmic step that despite changes to stack variables and the heap, the separation is preserved. The (\star) connective of separation logic, by contrast, allows explicit description of structure separation and the frame rule effortlessly exploits that separation.

When we are dealing with a single list or tree, aliasing doesn't matter, because it simply can't occur. Even in cases where there is substructure sharing – when two trees point to the same list, for example – changes to the substructure don't affect the structures that point to it. When we are dealing with a DAG or a graph via a global invariant, aliasing is straightforward to deal with once separation has been described. Following pointers and changing pointers, which is the name of the game in very many algorithms, is perfect separation logic territory. That's why we say that reasoning in separation logic embraces the dangling pointer: it doesn't matter what is attached to the end of a dangling pointer, because some other part of the program will separately preserve an invariant of that separate data structure.

A program that manipulates a cyclic graph, or even a directed acyclic graph, has to deal explicitly with the possibility that there can be multiple pointers to a substructure. In a cyclic graph it's obviously necessary to avoid infinite recursions or endless repetitions through the same sequence of vertices. Even in an acyclic graph it's normally a good idea to avoid considering the same subgraph many times. In effect, pointers are being used as surrogates for the structures they index, and the existence of other identical pointers – aliases – matters a great deal. Programs which deal with structures that can have internal sharing must keep some kind of record of substructures encountered, either by marking them as they are encountered, or by recording them in a list or a spanning tree (or whatever).

Our canonical problem is a recursive algorithm which copies a rooted, directed, ordered, possibly cyclic graph. In such a graph there can be several pointers to the same subgraph. In the copy there must be just the same number of pointers, in similar positions, to a similar substructure: that is, the aliasing relationships of the original have to be reproduced in the copy. Our specifications therefore have to describe aliasing. Because we want to specify the


algorithm in terms of its effect on a subgraph which may include aliased pointers, we need to specify some of the properties of the graph outside the footprint, but we want to keep that part of the specification to a minimum, to make the reasoning as local as possible.

Previous successful uses of separation logic in this area, notably Yang's treatment [13] of the Schorr-Waite graph-marking algorithm [10] and Torp-Smith's proof [11] of the Cheney two-space garbage-collection algorithm [3], have used global invariants to describe the total heap being manipulated by a program. A particular micro-action of the program, such as marking a heap cell or assigning a pointer, uses the frame rule to pick out and modify those parts of the invariant that are affected. Aliasing structures are described by describing sets of locations. In these problems a global invariant is appropriate: there is no recursion, the whole heap is available at every point of the program, arithmetic is sometimes necessary to distinguish one sort of pointer from another, and the program always needs to be able to access marks or pointers assigned at an earlier stage.

When dealing with recursive algorithms that deal with DAGs and graphs, it is clear that it isn't satisfactory to deal with global invariants. The algorithms are compact, they obviously operate on substructures, and it surely should be possible to reason locally in the same sort of way that we deal with list and tree algorithms. It's likely that the specifications will be significantly more complicated than with lists and trees, just because we have to describe aliasing. It's in that spirit that the work described below was undertaken. It's important to stress that if our reasoning is more 'local' than [13] and [11], it's because our problem algorithms are chosen to make it possible.

4 Specification style

In specifying programs, following Reynolds [8], we describe data structures by using a predicate which holds when a heap contains a representation of a data structure described by a formula. For example the tree



is described by the formula

$$\text{Node}(\text{Node}(\text{Tip } 1) (\text{Tip } 2)) \\ (\text{Node}(\text{Tip } 3) (\text{Node Empty} (\text{Tip } 4))) \quad (1)$$

The grammar of such descriptions, of possibly-empty trees with integers at their tips, is

$$T ::= \text{Node } T \ T \mid \text{Tip int} \mid \text{Empty}$$

A heap can represent a Node by a triple of 1 with two tree pointers, a Tip by a pair of 0 with the tip value, and Empty by nil. These decisions can be set out in a heap predicate

$$\text{tree nil Empty} \hat{=} \mathbf{emp} \\ \text{tree } t (\text{Tip } \alpha) \hat{=} t \mapsto 0, \alpha \\ \text{tree } t (\text{Node } \lambda \rho) \hat{=} \exists l, r. (t \mapsto 1, l, r \star \text{tree } l \lambda \star \text{tree } r \rho) \quad (2)$$

This predicate describes only the shape of a heap which represents a tree and not in general its domain (i.e. its position). There are usually many heaps which make a predicate formula true. For example,

the heap

$$2 \mapsto 0, 2 \star 6 \mapsto 1, \text{nil}, 80 \star 10 \mapsto 1, 266, 19 \star 19 \mapsto 1, 105, 6 \star 48 \mapsto 0, 1 \star 80 \mapsto 0, 4 \star 105 \mapsto 0, 3 \star 266 \mapsto 1, 48, 2$$

satisfies tree $10 E$ where E is the tree formula of (1). The same formula is also satisfied by

$$4 \mapsto 0, 1 \star 7 \mapsto 1, 4, 13 \star 10 \mapsto 1, 7, 100 \star 13 \mapsto 0, 2 \star 15 \mapsto 0, 3 \star 52 \mapsto 0, 4 \star 66 \mapsto 1, \text{nil}, 52 \star 100 \mapsto 1, 15, 66$$

It's then possible to specify simple recursive algorithms like *copytree*, which is given a pointer to a tree in the heap and produces a similarly shaped copy:

$$\{\text{tree } t \tau\} t' := \text{copytree } t \{\text{tree } t \tau \star \text{tree } t' \tau\}$$

Because the pre- and post-conditions specify the shape but not the domain of the heap, this specification doesn't require that the input heap is part of the output. A perverse *copytree* would be permitted to throw away part or all of its input and even assign to t , provided only that it leaves two separate τ -shaped heaps on exit. For the rest of this paper we ignore that kind of problem, which can be circumvented in ways that would add no illumination to our discussion.

5 Inclusion sharing

Our first example of internal structure sharing is the inclusion of one structure within another: a list within a tree. The idiom of the solution is widely useful, and we exploit it in our treatment of cyclic graphs with forwarding pointers below.

Consider a binary tree with values at its tips and no null subtrees. The grammar of formulae describing such trees is

$$B ::= \text{Node } B B \mid \text{Tip val}$$

The fringe of such a tree is a sequence of the values at its tips, which can be computed from a tree formula ($++$ is sequence append):

$$\begin{aligned} \text{fringe } (\text{Tip } v) &\triangleq \langle v \rangle \\ \text{fringe } (\text{Node } \lambda \rho) &\triangleq \text{fringe } \lambda ++ \text{fringe } \rho \end{aligned} \quad (3)$$

Richard Bird posed us the problem of deriving ‘without fuss’ a program which links together the fringe of this kind of tree. He observed that if nodes and tips are both represented in the heap by a triple (l, v, r) , tips would naturally have two nil subtree pointers l and r , either of which would be enough to distinguish them from nodes, and the other pointer slot could then be exploited to link together the tips. Nodes would have an unused value slot. That corresponds to the heap predicate

$$\begin{aligned} \text{btree } t (\text{Tip } v) &\triangleq t \mapsto \text{nil}, v, - \\ \text{btree } t (\text{Node } \lambda \rho) &\triangleq \exists l, r. (t \mapsto l, -, r \star \text{btree } l \lambda \star \text{btree } r \rho) \end{aligned} \quad (4)$$

A list is conventionally represented by a linked sequence of (value, pointer) records, the last of which contains a nil pointer. In the heap representation of a *Tip*, the second two slots can be viewed as a component of a list. Rather than defining a nil-terminated list, we need to define a list segment – a sequence of records the last of which contains a continuation pointer y , the whole representing a sequence vs .

$$\begin{aligned} \text{lseg } y y \langle \rangle &\triangleq \text{emp} \\ \text{lseg } x y \langle (v) ++ vs \rangle &\triangleq \exists x'. (x \mapsto v, x' \star \text{lseg } x' y vs) \end{aligned} \quad (5)$$


Note that a segment is a sequence of separated records, but the continuation pointer y is unconstrained. In our examples y will always point to another segment or be nil, but we don't have to rely on that when checking whether a heap is an *lseg*.

5.1 Algorithm and specification

Since our business, at this stage of our understanding of separation logic, is post-hoc verification rather than derivation, we focussed on what seemed to us to be the obvious algorithm:

$$\begin{aligned} \text{fringelink } t c &\triangleq \text{if } [t] = \text{nil} \text{ then} \\ &\quad [t+2] := c; t+1 \\ &\quad \text{else} \\ &\quad \quad \text{fringelink } [t] (\text{fringelink } [t+2] c) \\ &\quad \text{fi} \end{aligned} \quad (6)$$

The continuation pointer argument c links the fringes of adjacent subtrees together, and overall *fringelink* t nil returns a pointer to the nil-terminated fringe list of tree t . The input footprint of *fringelink* $t c$ is clearly the tree pointed to by t , and the result points to a list taking up part of the tree space:



A heap which was both fringe F and tree T would be $F \wedge T$. The fringe is normally smaller than the tree, so the output is can be described as $(F \star \text{True}) \wedge T$. The specification of the algorithm is then

$$\begin{aligned} \{\text{btree } t \tau\} \\ \text{res} := \text{fringelink } t c \\ \{\text{lseg } \text{res } c (\text{fringe } \tau) \star \text{True}\} \wedge \text{btree } t \tau \end{aligned} \quad (7)$$

5.2 Proof

To clarify the proof it's helpful to compile the nested function call of the else arm into a sequence of assignments. The proof of the modified algorithm must then show

$$\begin{aligned} \{\text{btree } t \tau\} \\ \text{if } [t] = \text{nil} \text{ then} \\ \quad \{\text{btree } t (\text{Tip } v)\} \\ \quad [t+2] := c; t+1 \\ \quad \{(\text{lseg } (t+1) c \langle v \rangle \star \text{True}) \wedge \text{btree } t (\text{Tip } v)\} \\ \quad \text{else} \\ \quad \quad \{\text{btree } t (\text{Node } \lambda \rho)\} \\ \quad \quad \text{res1} := \text{fringelink } [t+2] c; \\ \quad \quad \text{res2} := \text{fringelink } [t] \text{res1}; \\ \quad \quad \text{res2} \\ \quad \quad \{(\text{lseg } \text{res2 } c (\text{fringe } \lambda ++ \text{fringe } \rho) \star \text{True}) \wedge \\ \quad \quad \quad \{\text{btree } t (\text{Node } \lambda \rho)\} \\ \quad \quad \text{fi} \\ \quad \{(\text{lseg } \text{res } c (\text{fringe } \tau) \star \text{True}) \wedge \text{btree } t \tau\} \end{aligned}$$

The frame rule is used several times in the proof itself, and we've introduced a notational device to clarify those steps. The deduction

$$\frac{\{Q\}C\{R\}}{\{P \star Q\}C\{P \star R\}}$$

else
 $\{ \text{btree } t \tau \wedge [t] \neq \text{nil} \} \cdot \{ \text{btree } t (\text{Node } \lambda \rho) \} \cdot \{ t \mapsto l, _, r \star \text{btree } l \lambda \star \text{btree } r \rho \}$
 Framed: $\langle \{ \text{btree } r \rho \} \text{res1} := \text{fringelink } [t + 2] \text{ c}; \{ (\text{lseg } \text{res1 } c (\text{fringe } \rho) \star \text{True}) \wedge \text{btree } r \rho \} \rangle$
 $\{ t \mapsto l, _, r \star \text{btree } l \lambda \star ((\text{lseg } \text{res1 } c (\text{fringe } \rho) \star \text{True}) \wedge \text{btree } r \rho) \}$
 Framed: $\langle \{ \text{btree } l \lambda \} \text{res2} := \text{fringelink } [t] \text{ res1}; \{ (\text{lseg } \text{res2 } \text{res1} (\text{fringe } \lambda) \star \text{True}) \wedge \text{btree } l \lambda \} \rangle$
 $\{ t \mapsto l, _, r \star ((\text{lseg } \text{res2 } \text{res1} (\text{fringe } \lambda) \star \text{True}) \wedge \text{btree } l \lambda) \star ((\text{lseg } \text{res1 } c (\text{fringe } \rho) \star \text{True}) \wedge \text{btree } r \rho) \} \cdot$
 $\{ (\text{lseg } \text{res2 } \text{res1} (\text{fringe } \lambda) \star \text{lseg } \text{res1 } c (\text{fringe } \rho) \star \text{True}) \wedge (t \mapsto l, _, r \star \text{btree } l \lambda \star \text{btree } r \rho) \} \cdot$
 $\{ (\text{lseg } \text{res2 } c (\text{fringe } \lambda ++ \text{fringe } \rho) \star \text{True}) \wedge \text{btree } t (\text{Node } \lambda \rho) \}$
 res2
 $\{ (\text{lseg } \text{res } c (\text{fringe } \lambda ++ \text{fringe } \rho) \star \text{True}) \wedge \text{btree } t (\text{Node } \lambda \rho) \}$

Figure 1. fringelink proof, Node case

$\text{pdag nil Empty } U U \hat{=} \text{emp}$
 $\text{pdag } d (\text{Ptr } x) U U \hat{=} U x = d \wedge \text{emp}$
 $\text{pdag } d (x : \text{Tip } \alpha) U V \hat{=} x \notin \text{dom } U \wedge V = U \oplus (x : d) \wedge d \mapsto 0, \alpha$
 $\text{pdag } d (x : \text{Node } \lambda \rho) U V \hat{=} \exists l, r, U', V'. (x \notin \text{dom } U \wedge x \notin \text{dom } V' \wedge V = V' \oplus (x : d) \wedge d \mapsto 1, l, r \star \text{pdag } l \lambda U U' \star \text{pdag } r \rho U' V')$

Figure 2. Heap predicate for partial DAGs

$\{ \text{pdag } d \delta U V \wedge \text{ran } U = \text{dom } f \wedge \text{nil} \notin \text{ran } U \wedge d \neq \text{nil} \wedge d \notin \text{dom } f \}$
 elsif $d.\text{tag} = 0$ then
 $\{ \text{pdag } d (x : \text{Tip } \alpha) U (U \oplus (x : d)) \wedge \text{ran } U = \text{dom } f \wedge \text{nil} \notin \text{ran } U \} \cdot$
 $\{ d \mapsto 0, \alpha \wedge x \notin \text{dom } U \wedge \text{ran } U = \text{dom } f \wedge \text{nil} \notin \text{ran } U \}$
 $d' := \text{new}(0, d.\text{val});$
 $\{ d \mapsto 0, \alpha \star d' \mapsto 0, \alpha \wedge x \notin \text{dom } U \wedge x \notin \text{dom } (f \bullet U) \wedge \text{ran } U = \text{dom } f \wedge \text{nil} \notin \text{ran } U \} \cdot$
 $\{ \text{pdag } d (x : \text{Tip } \alpha) U (U \oplus (x : d)) \star \text{pdag } d' (x : \text{Tip } \alpha) (f \bullet U) ((f \bullet U) \oplus (x : d')) \wedge \text{ran } U = \text{dom } f \wedge \text{nil} \notin \text{ran } U \}$
 $d', f \oplus (d : d')$
 $\{ (\text{pdag } d (x : \text{Tip } \alpha) U (U \oplus (x : d)) \wedge (d \mapsto _ \star \text{True})) \star \text{pdag } d (x : \text{Tip } \alpha) (f \bullet U) ((f \oplus (d : d')) \bullet (U \oplus (x : d))) \wedge$
 $\{ \text{ran}(U \oplus (x : d)) = \text{dom}(f \oplus (d : d')) \wedge \text{nil} \notin \text{ran}(U \oplus (x : d)) \}$

Figure 3. Slow copydag, Tip case

else
 $\{ \text{pdag } d \delta U V \wedge \text{ran } U = \text{dom } f \wedge \text{nil} \notin \text{ran } U \wedge (\forall x \in \text{ran } U \cdot x \not\mapsto _) \wedge d \neq \text{nil} \wedge d \notin \text{dom } f \wedge d.\text{tag} \neq 0 \} \cdot$
 $\{ \text{pdag } d (x : \text{Node } \lambda \rho) U V \wedge \text{ran } U = \text{dom } f \wedge \text{nil} \notin \text{ran } U \wedge (\forall x \in \text{ran } U \cdot x \not\mapsto _) \} \cdot$
 $\{ d \mapsto 1, L, R \star \text{pdag } L \lambda U U' \star \text{pdag } R \rho U' V' \wedge \text{ran } U = \text{dom } f \wedge \text{nil} \notin \text{ran } U \wedge (\forall x \in \text{ran } U \cdot x \not\mapsto _) \wedge x \notin \text{dom } U \wedge x \notin \text{dom } V' \}$
 Framed: $\left\langle \begin{array}{l} l, f' := \text{copydag } d.\text{left } f \\ \{ \text{pdag } L \lambda U U' \star \text{pdag } l \lambda (f \bullet U) (f' \bullet U') \wedge \text{ran } U' = \text{dom } f' \wedge \text{nil} \notin \text{ran } U' \} \end{array} \right\rangle;$
 $\left\langle \begin{array}{l} (d \mapsto 1, L, R \star (\text{pdag } L \lambda U U' \wedge ((\forall \ast y \in \text{ran}(U' - U) \cdot y \mapsto _) \star \text{True})) \star \text{pdag } R \rho U' V' \wedge (\forall x \in \text{ran } U \cdot x \not\mapsto _) \star \\ \{ \text{pdag } l \lambda (f \bullet U) (f' \bullet U') \wedge \text{ran } U' = \text{dom } f' \wedge \text{nil} \notin \text{ran } U' \wedge x \notin \text{dom } U \wedge x \notin \text{dom } V' \} \end{array} \right\rangle$
 Framed: $\left\langle \begin{array}{l} \{ (\text{pdag } R \rho U' V' \wedge \text{ran } U' = \text{dom } f' \wedge \text{nil} \notin \text{ran } U' \wedge (\forall x \in \text{ran } U' \cdot x \not\mapsto _)) \\ r, f'' := \text{copydag } d.\text{right } f' \\ \{ \text{pdag } R \rho U' V' \star \text{pdag } r \rho (f' \bullet U') (f'' \bullet V') \wedge \text{ran } V' = \text{dom } f'' \wedge \text{nil} \notin \text{ran } V' \} \end{array} \right\rangle;$
 $\left\langle \begin{array}{l} d \mapsto 1, L, R \star \text{pdag } L \lambda U U' \star \text{pdag } R \rho U' V' \star \text{pdag } l \lambda (f \bullet U) (f' \bullet U') \star \text{pdag } r \rho (f' \bullet U') (f'' \bullet V') \wedge \\ \{ \text{ran } V' = \text{dom } f'' \wedge \text{nil} \notin \text{ran } V' \wedge x \notin \text{dom } U \wedge x \notin \text{dom } V' \} \end{array} \right\rangle$
 $d' := \text{new}(1, l, r);$
 $\left\langle \begin{array}{l} d \mapsto 1, L, R \star \text{pdag } L \lambda U U' \star \text{pdag } R \rho U' V' \star d' \mapsto 1, l, r \star \text{pdag } l \lambda (f \bullet U) (f' \bullet U') \star \text{pdag } r \rho (f' \bullet U') (f'' \bullet V') \wedge \\ \{ \text{ran } V' = \text{dom } f'' \wedge \text{nil} \notin \text{ran } V' \wedge x \notin \text{dom } U \wedge x \notin \text{dom } V' \} \end{array} \right\rangle$
 $d', f'' \oplus (d, d')$
 $\{ \text{pdag } d (x : \text{Node } \lambda \rho) U (V' \oplus (x : d)) \star \text{pdag } d' (x : \text{Node } \lambda \rho) (f \bullet U) ((f'' \oplus (d, d')) \bullet (V' \oplus (x : d))) \wedge$
 $\{ \text{ran}(V' \oplus (x : d)) = \text{dom}(f'' \oplus (d : d')) \wedge \text{nil} \notin \text{ran}(V' \oplus (x : d)) \}$

Figure 4. Slow copydag, Node case

$\text{graph nil Empty } U U \hat{=} \text{emp}$
 $\text{graph } g (\text{Ptr } x) U U \hat{=} U x = g \wedge \text{emp}$
 $\text{graph } g (x : \text{GNode } \alpha \lambda \rho) U V \hat{=} \exists l, r, U'. (x \notin \text{dom } U \wedge g \mapsto \alpha, l, r \star \text{graph } l \lambda (U \oplus (x : g)) U' \star \text{graph } r \rho U' V)$

Figure 5. Heap predicate for partial graphs

$\{ (\text{fwdgraph } g \gamma U V \wedge ((\forall \ast x \in \text{ran}(V - U) \cdot x \mapsto \text{nil}) \star \text{True})) \star (\forall \ast y \in \text{ran } U \cdot y \mapsto f y) \wedge \text{dom } f = \text{ran } U \wedge \text{nil} \notin \text{ran } f \}$
 $g' := \text{copyfwdgraph } g$
 $\left\langle \begin{array}{l} \exists f'. \left((\text{fwdgraph } g \gamma U V \wedge ((\forall \ast x \in \text{ran}(V - U) \cdot x \mapsto f' x) \star \text{True})) \star \text{fwdgraph } g' \gamma (f \bullet U) ((f \oplus f') \bullet V) \star \right) \\ \left((\forall \ast y \in \text{ran } U \cdot y \mapsto f y) \wedge \text{dom}(f \oplus f') = \text{ran } V \wedge \text{nil} \notin \text{ran}(f \oplus f') \right) \end{array} \right\rangle$

Figure 6. Specification of fast copygraph

is shown linearly as

$$\begin{array}{l} \{P \star Q\} \\ \text{Framed: } \langle \{Q\} C \{R\} \rangle \\ \{P \star R\} \end{array}$$

We have taken some shortcuts in the presentation of the framed command C in this and other examples where arguments to recursive calls are not pure expressions. *fringelink* $[t]$ *resl*, for example, mentions the cell pointed to by t but we don't include that in the frame in figure 1. De-sugaring of procedure calls would make the command $z := [t]; \text{fringelink } z \text{ resl}$, and it's clear that the framed procedure call doesn't need access to the cell.

The Tip case constructs a singleton list. The only remarkable steps are the introduction of \wedge and the replacement of $t \mapsto \text{nil}$ by True:

$$\begin{array}{l} \text{if } [t] = \text{nil} \text{ then} \\ \quad \{\text{btree } t \wedge [t] = \text{nil}\} \therefore \{\text{btree } t (\text{Tip } v)\} \therefore \\ \quad \{t \mapsto \text{nil}, v, -\} \\ \quad \text{Framed: } \langle \{t + 2 \mapsto -\} [t + 2] := c; \{t + 2 \mapsto c\} \rangle \\ \quad \{t \mapsto \text{nil}, v, c\} \therefore \\ \quad \{t \mapsto \text{nil}, v, c \wedge t \mapsto \text{nil}, v, c\} \therefore \\ \quad \{(t + 1 \mapsto v, c \star t \mapsto \text{nil}) \wedge t \mapsto \text{nil}, v, c\} \therefore \\ \quad \{(\text{lseg } (t + 1) c \langle v \rangle \star \text{True}) \wedge \text{btree } t (\text{Tip } v)\} \\ \quad t + 1 \\ \quad \{(\text{lseg } \text{res } c \langle v \rangle \star \text{True}) \wedge \text{btree } t (\text{Tip } v)\} \end{array} \quad (8)$$

The Node case (figure 1) is not quite so straightforward. After the second framed step, the postcondition is of the form $A \star ((B \star \text{True}) \wedge C) \star ((D \star \text{True}) \wedge E)$. This can be rearranged, using the rules of BI plus the fact that $\text{True} \star \text{True} = \text{True}$, into $(B \star D \star \text{True}) \wedge (A \star C \star E)$. It's possible to show by induction on α that $\text{lseg } x x' \alpha \star \text{lseg } x' y \beta \vdash \text{lseg } x y (\alpha ++ \beta)$, which justifies the final conclusion.

6 DAGs and graphs

Cyclic graphs are a major challenge to separation logic. The heap representation of a graph involves, in principle, a tangle of pointers between interpenetrating structures. Records aren't labelled to say that they are shared. Cycles of edges aren't marked.

Despite that, graph algorithms are remarkably simple. A few basic tricks – e.g. marking, forwarding pointers – appear to cover most eventualities. It surely ought to be possible to deal with the problem as simply as the algorithms do.

DAGs – directed acyclic graphs – are halfway between trees and graphs: substructures are shared but there are no cycles. The essential difference between a DAG and a graph is that the root of a DAG is separate from its children. In search of simple solutions, we begin by considering the problem of copying a DAG, producing a copy which has exactly the same aliasing relationships as the original. As with the *copytree* example above, we ignore the necessity to state that the input DAG is undisturbed.

Although it would be possible to justify an algorithm by rational reconstruction from a specification, in truth we started with a program. One way – the slow way – to copy a DAG is to scan the structure, building up an association list of (original pointer, copy pointer) pairs as you go. Each time you come across a substructure, you check to see if it's recorded in the list: if it is, you use the already-recorded copy pointer; if it's not, you copy it and record a new pair in the list. Abstracting from an association list to a forwarding function, we have the following procedure, which takes

as argument a pointer d to a binary DAG,¹ represented in the heap very like the binary trees of (2), and a forwarding function f representing copies made so far. It returns a pointer to a copy of the DAG, plus an updated forwarding function.


$$\begin{array}{l} \text{copydag } d f \triangleq \text{if } d = \text{nil} \text{ then nil, } f \\ \quad \text{elsif } d \in \text{dom } f \text{ then } f d, f \\ \quad \text{elsif } d.\text{tag} = 0 \text{ then} \\ \quad \quad d' := \text{new}(0, d.\text{val}); d', f \oplus (d : d') \\ \quad \text{else} \\ \quad \quad l, f' := \text{copydag } d.\text{left } f; \\ \quad \quad r, f'' := \text{copydag } d.\text{right } f'; \\ \quad \quad d' := \text{new}(1, l, r); \\ \quad \quad d', f'' \oplus (d : d') \\ \text{fi} \end{array} \quad (9)$$

We describe this below as an aliasing-discovery copying algorithm. Line 2 deals with aliasing; the rest of the algorithm is just tree copying except that in the else case the recursive calls augment the forwarding function f – the aliasing description – as they go. Substructures in a DAG can't share with the root, and to emphasise the point we copy the root after dealing with its children.

The ideal specification of this algorithm would be that given a pointer to a closed DAG – one with no dangling pointers – and a null forwarding function, it produces a heap which contains the original DAG and a separate identically-aliased copy, with a forwarding function which describes an isomorphism between original and copy. We approach that ideal in stages.

6.1 Describing a DAG

To specify and verify (9) we must first be able to give a syntactic description of a DAG. Our binary DAGs, like binary trees, can contain nodes and tips. We allow null subDAGs. Binding formulae 'let $x = D$ in D' ', with Ptr x references to defined names, are an obvious way to describe sharing of structures. For example, the DAG



is described by the formula

$$\text{let } c = \text{Tip } 17 \text{ in Node } (\text{Node Empty } (\text{Ptr } c)) \quad (\text{Node } (\text{Ptr } c) \text{ Empty}) \quad (10)$$

The grammar of this kind of DAG formula is

$$D ::= \text{Empty} \mid \text{Tip int} \mid \text{Node } D D \mid \text{Ptr var} \mid \text{let var} = D \text{ in } D \quad (11)$$

The corresponding heap predicate takes a pointer d and a formula δ as arguments, just as (2) and (4) do. The third argument U is a mapping from bound names to dangling pointers which index shared structures in the heap – in effect, continuation pointers for each of the undefined names in δ . Unreferenced definitions would make disconnected structures in the heap, so we impose a relevance con-

¹It wouldn't be difficult to generalise to n -ary DAGs, but it wouldn't add anything to the discussion.

dition on binding formulae.

$$\begin{aligned}
&\text{lidag nil Empty } U \triangleq \mathbf{emp} \\
&\text{lidag } d \text{ (Tip } \alpha) U \triangleq d \mapsto 0, \alpha \\
&\text{lidag } d \text{ (Node } \lambda \rho) U \triangleq \exists l, r. \begin{pmatrix} d \mapsto 1, l, r \star \text{lidag } l \lambda U \star \\ \text{lidag } r \rho U \end{pmatrix} \\
&\text{lidag } d \text{ (Ptr } x) U \triangleq U x = d \wedge \mathbf{emp} \\
&\text{lidag } d \text{ (let } x = \delta' \text{ in } \delta) U \triangleq \exists d'. \begin{pmatrix} x \text{ free in } \delta \wedge \text{lidag } d' \delta' U \star \\ \text{lidag } d \delta (U \oplus (x : d')) \end{pmatrix}
\end{aligned} \tag{12}$$

7 Partial graphs

Unfortunately the *lidag* predicate doesn't suit our purposes. In the heap there is no advance signalling of internal sharing, nothing to correspond to a let-in construction. The algorithm of (9) discovers internal sharing by noting substructures as it comes across them and recognising when the same substructure is referenced again. It would be possible to construct a two-stage copy process, with a first pass which discovered the aliasing structure and a second, based on the *lidag* predicate, which did the copying. But then the first pass would have to be based on a heap predicate which doesn't assume knowledge of aliasing, which is just the problem that we started with. It's obvious that it's no harder to verify an aliasing-discovery copying algorithm like (9), which discovers aliasing as it goes, than to take the two-stage approach.

It's clear that any adequate DAG copying algorithm must in effect discover a spanning tree of its argument, and that (9) finds a spanning tree by treating the first occurrence of a subDAG encountered in a left-to-right scan as the defining instance, cutting links to the same structure which are encountered later. It's possible to transform a let-in formula like (10) into one which suits that particular biased search. Take the innermost let-in formula: push it inwards, preserving the relevance condition (D in $\text{let } x = D' \text{ in } D$ must contain $\text{Ptr } x$), and pushing into the left subDAG of a *Node* if there is a choice. The let binding will necessarily come to rest against the leftmost occurrence of a reference to its name. Replace $\text{let } x = D$ in $\text{Ptr } x$ with $x : D$. Repeat with other let bindings until all are removed. The formula of (10), for example, is transformed into


$$\text{Node (Node Empty } (c : \text{Tip } 17)) \text{ (Node (Ptr } c) \text{ Empty)}$$

We can easily read this formula left-to-right: $x : D$ labels a structure on first occurrence, and $\text{Ptr } x$ later refers to it.

Then observe that when a program scans the heap, it can't tell whether or not a substructure it has come across for the first time will be referenced later, so label *every* element, giving

$$\begin{aligned}
&a : \text{Node } (b : \text{Node Empty } (c : \text{Tip } 17)) \\
&\quad (d : \text{Node (Ptr } c) \text{ Empty})
\end{aligned}$$

This formula describes the fully-labelled DAG



The grammar of these formulae is

$$D ::= \text{Empty} \mid \text{Ptr lab} \mid \text{lab} : \text{Tip int} \mid \text{lab} : \text{Node } DD \tag{13}$$

These formulae describe *partial graphs*: graphs consisting of named vertices in which not every name maps to a successor set, because there is no requirement that $x : D$ must occur to define $\text{Ptr } x$. They are extremely similar to the structures described by Cardelli, Gardner and Ghelli in [2] though the motivation is different: their trees are given and the pointers defined as secondary; ours must be discovered, and the direction of a scan determines which pointers are part of the tree and which are aliases.

We stumbled across partial graphs as a means of describing the heap footprint of (9), plucking them out of the air as a transformation of the let-in formulae of (11), but they do appear to be semantically interesting in their own right. They suit graphs even better than they do DAGs, as we shall see. We haven't yet had time to investigate their properties. They may be awkward to work with – it's not quite so easy to say that a partial graph is closed, for example, as it is with a let-in formula, and it's necessary to impose a unique-labelling restriction – but they do suit our purposes rather well.

8 Describing partial DAGs

The heap predicate for DAGs based on the partial graph formulae of (13) is shown in figure 2. The algorithm (9) accumulates a forwarding function; the predicate describes a left-to-right evaluation that similarly accumulates an environment. By contrast, the environment of (12) is built up recursively. U is the environment mapping names which are not defined in the partial DAG (providing a collection of continuation pointers); V is the output environment which extends U to include also the names that are defined in the partial DAG. In the $x : \text{Node } \dots$ case, since we are dealing with a DAG, there mustn't be occurrences of $\text{Ptr } x$ within the node. It would be sufficient to say either that x isn't mapped by environment U or that it isn't mapped by V' , but we explicitly exclude it from both as a proof convenience. The unique labelling condition is imposed by the same restriction.

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.

There is a more subtle complication, however. In the *pdag* predicate of figure 2 the domain of the input environment U (names of allocated subDAGs) is constrained, but the range (pointers to corresponding heaps) is not. U must point somewhere, and in practice – in 'normal' elaboration starting with a closed formula and an empty environment – it points to already-allocated structures elsewhere in the heap. But we don't yet know how to say that in separation logic without dragging a global environment into our specifications. Like the continuation argument c in the *fringelink* algorithm (6), if we constrain $\text{ran } U$ too much our proofs become too complicated.

The consequence is that *pdag* (like *lidag*, which doesn't constrain $\text{ran } U$ either) doesn't describe exclusively DAG-like heaps. If δ is a closed DAG formula and *null* is the empty environment, then certainly $\text{pdag } d \delta \text{ null } V$ describes a DAG-like heap no matter what the values of d and δ . But it is easy to construct an environment and a formula which describe a cyclic heap structure: for example

$$\begin{aligned}
&\delta = a : (\text{Node (Ptr } b) \text{ Empty}) \\
&d = 10 \\
&U = (b : 10) \\
&V = (b : 10; a : 10)
\end{aligned} \tag{14}$$

$$\text{heap: } 10 \mapsto 1, 10, \text{nil}$$

is distinctly unDAG-like.²

Even if $\text{ran } U$ were defined to point outside the heap of $\text{pdag } d \delta U V$ it wouldn't be enough to ensure genuinely DAG-like heaps, because dangling pointers can by accident coincide with a pointer returned by new. We can't be sure that $x := \text{new}(1, d, d)$ will give us a non-cyclic structure, for example, if x happens to receive a value which is in $\text{ran } U$. Of course in 'normal' elaboration that can't happen: $\text{ran } U$ will always point to allocated structures, and new will never return a pointer to an already-allocated structure. It seems in practice that with our current understanding if we can't say the right thing it is better to say nothing. The potentially spurious heap predicate of figure 2 will suit us very well.

Although $\text{ran } U$ is unconstrained, $\text{ran}(V - U)$, the part of the output environment V that points to structures contained in $\text{pdag } d \delta U V$, is not. $(V - U)$ necessarily contains non-dangling pointers. It's easy to see that no matter what the value of U , we have

$$\text{pdag } d \delta U V \rightarrow (\text{pdag } d \delta U V \wedge ((\forall x, y \in \text{ran}(V - U) \cdot y \mapsto _) \star \text{True})) \quad (15)$$

($\forall \star$ is the iterated conjunction introduced in [11].)

8.1 Copying a partial DAG slowly

If we disingenuously imagine that the representation of a forwarding function takes no heap space,³ the function of (9) takes a DAG-like heap and produces a similarly-shaped copy, leaving the original unchanged. The resulting specification, where $f \bullet g$ is function composition – i.e. $(f \bullet g) x = f(g x)$, is

$$\left\{ \begin{array}{l} \text{pdag } d \delta U V \wedge (\forall x \in \text{ran } U \cdot x \not\mapsto _) \wedge \\ \text{ran } U = \text{dom } f \wedge \text{nil} \notin \text{ran } U \\ d', f' := \text{copydag } d f \\ \text{pdag } d \delta U V \star \text{pdag } d' \delta (f \bullet U) (f' \bullet V) \wedge \\ \text{ran } V = \text{dom } f' \wedge \text{nil} \notin \text{ran } V \end{array} \right\} \quad (16)$$

The specification requires that the input environment U contains only dangling pointers, because the algorithm uses the test $d \in \text{dom } f$ to distinguish them from others which point into the heap. But in the postcondition we can't easily require that $\text{ran}(f \bullet U)$ is similarly dangling: once again, in 'normal' elaboration it must be so, but without incorporating a global invariant we don't know how to specify it. The laxity of the pdag predicate allows us to say nothing about $\text{ran}(f \bullet U)$ and to allow the constraints of composition of substructures to provide the necessary information, just like the treatment of the c argument in figure 1. That property of $f \bullet U$ is in any case irrelevant overall, since (16), if established, implies

$$\left\{ \begin{array}{l} \text{pdag } d \delta \text{null } V \\ d', f' := \text{copydag } d \text{null} \\ \text{pdag } d \delta \text{null } V \star \text{pdag } d' \delta \text{null } (f' \bullet V) \end{array} \right\} \quad (17)$$

We show a proof in stages. The Empty case exploits the fact that

²We are grateful to the anonymous reviewer who pointed out this counter-example to us.

³This is merely a simplifying assumption. Being realistic about association lists would add nothing to the discussion, since the values in the list are only compared and copied, never dereferenced.

$\text{nil} \notin \text{ran } U$ but is otherwise trivial:

$$\left\{ \begin{array}{l} \text{pdag } d \delta U V \wedge \text{ran } U = \text{dom } f \wedge \text{nil} \notin \text{ran } U \\ \text{if } d = \text{nil} \text{ then} \\ \quad \left\{ \begin{array}{l} \text{pdag } \text{nil} \delta U V \wedge \text{ran } U = \text{dom } f \wedge \text{nil} \notin \text{ran } U \end{array} \right\} \cdot: \\ \quad \left\{ \text{pdag } \text{nil} \text{Empty } U U \wedge \text{ran } U = \text{dom } f \wedge \text{nil} \notin \text{ran } U \end{array} \right\} \cdot: \\ \quad \left\{ \text{emp} \wedge \text{ran } U = \text{dom } f \wedge \text{nil} \notin \text{ran } U \end{array} \right\} \cdot: \\ \quad \left\{ \text{emp} \star \text{emp} \wedge \text{ran } U = \text{dom } f \wedge \text{nil} \notin \text{ran } U \right\} \\ \quad \text{nil}, f \\ \quad \left\{ \begin{array}{l} \text{pdag } \text{nil} \text{Empty } U U \star \text{pdag } \text{nil} \text{Empty } (f \bullet U) (f \bullet U) \wedge \\ \text{ran } U = \text{dom } f \wedge \text{nil} \notin \text{ran } U \end{array} \right\} \end{array} \right\} \quad (18)$$

In the Ptr case we need to argue

$$d \in \text{dom } f \wedge \text{dom } f = \text{ran } U \rightarrow \exists x \cdot U x = d$$

and then because $\text{ran } U$ doesn't point at anything, we know we are dealing with Ptr x .

$$\left\{ \begin{array}{l} \text{pdag } d \delta U V \wedge \text{ran } U = \text{dom } f \wedge \text{nil} \notin \text{ran } U \wedge \\ (\forall x \in \text{ran } U \cdot x \not\mapsto _) \wedge d \neq \text{nil} \\ \text{elsif } d \in \text{dom } f \text{ then} \\ \quad \left\{ \begin{array}{l} \text{pdag } d \delta U V \wedge \text{ran } U = \text{dom } f \wedge (\forall x \in \text{ran } U \cdot x \not\mapsto _) \wedge \\ \text{nil} \notin \text{ran } U \wedge d \neq \text{nil} \wedge d \in \text{ran } U \end{array} \right\} \cdot: \\ \quad \left\{ \begin{array}{l} \text{pdag } d \delta U V \wedge \text{ran } U = \text{dom } f \wedge d \not\mapsto _ \wedge \\ \text{nil} \notin \text{ran } U \wedge d \neq \text{nil} \wedge U x = d \end{array} \right\} \cdot: \\ \quad \left\{ \text{pdag } d (\text{Ptr } x) U U \wedge \text{ran } U = \text{dom } f \wedge \text{nil} \notin \text{ran } U \wedge U x = d \right\} \cdot: \\ \quad \left\{ \text{emp} \wedge \text{ran } U = \text{dom } f \wedge \text{nil} \notin \text{ran } U \wedge U x = d \right\} \cdot: \\ \quad \left\{ \text{emp} \star \text{emp} \wedge \text{ran } U = \text{dom } f \wedge \text{nil} \notin \text{ran } U \wedge U x = d \right\} \\ \quad f d, f \\ \quad \left\{ \begin{array}{l} \text{pdag } d (\text{Ptr } x) U U \star \text{pdag } (f d) (\text{Ptr } x) (f \bullet U) (f \bullet U) \wedge \\ \text{ran } U = \text{dom } f \wedge \text{nil} \notin \text{ran } U \end{array} \right\} \end{array} \right\} \quad (19)$$

The Tip case (figure 3) is straightforward. The Node case (figure 4) uses the input property $(\forall x \in \text{ran } U \cdot x \not\mapsto _)$ in the first recursive call. Before the second call we note that $\text{ran}(U' - U)$ points into $\text{pdag } L \lambda U U'$, which is separated by a (*) from $\text{pdag } R \rho U' V'$. It follows that $\text{ran}(U' - U)$ dangles outside $\text{pdag } R \rho U' V'$; since $\text{ran } U$ dangles outside as well, we know that the whole of $\text{ran } U'$ dangles outside, which gives the necessary precondition for the second recursive call.

8.2 Disposing of a partial DAG slowly

Trees can be disposed recursively: subtree, subtree, ..., root. DAGs need more care, because we mustn't follow a pointer to a shared but already-disposed structure. The development so far is enough to support an algorithm for disposing DAGs. Rather than a forwarding function we keep a single list ds of disposed nodes.

$$\begin{array}{l} \text{disposedag } d \text{ ds} \hat{=} \text{if } d = \text{nil} \vee d \in \text{ds} \text{ then } ds \\ \quad \text{elsif } d.\text{tag} = 0 \text{ then} \\ \quad \quad \text{disposed } 2; d :: ds \\ \quad \text{else} \\ \quad \quad ds' := \text{disposedag } d.\text{left } ds; \\ \quad \quad ds'' := \text{disposedag } d.\text{right } ds'; \\ \quad \quad \text{disposed } 3; \\ \quad \quad d :: ds'' \\ \quad \text{fi} \end{array} \quad (20)$$

The specification must once again require that the input DAG is properly DAG-shaped. This time we don't pretend to ignore the space taken up by the list, so to get rid of a DAG you need $\text{disposelist}(\text{disposedag } d \text{ nil})$.

$$\begin{aligned} & \{(pdag\ d\ \delta\ U\ V \wedge (\forall x \in \text{ran}\ U \cdot x \not\leftrightarrow _) \star \text{lseg}\ ds\ c\ (\text{ran}\ U))\} \\ & \quad ds' := \text{disposedag}\ d\ ds \\ & \{ \text{lseg}\ ds' \ c\ (\text{ran}\ V) \} \end{aligned} \quad (21)$$

The proof is extremely similar to the proof of *copydag* and is omitted.

8.3 Copying and disposing a partial graph, slowly

It's as easy to copy a cyclic graph as it is to copy a DAG. In a graph we put values at nodes and there is no need for tips

$$G ::= \text{Empty} \mid \text{Ptr var} \mid \text{var} : \text{GNode int } G\ G \quad (22)$$

and use a representation predicate (figure 5) which allows for the fact that we can have cycles. The graph-copying algorithm is then a minor modification of *copydag*, extending the forwarding function before copying subgraphs in order to deal with cycles.

$$\begin{aligned} \text{copygraph } g\ f \hat{=} & \text{ if } g = \text{nil} \text{ then nil, } f \\ & \text{ elsef } g \in \text{dom } f \text{ then } f\ g, f \\ & \text{ else} \\ & \quad g' := \text{new}(g.\text{val}, \text{nil}, \text{nil}); \\ & \quad g'.\text{left}, f' := \text{copygraph } g.\text{left } (f \oplus (g : g')); \\ & \quad g'.\text{right}, f'' := \text{copygraph } g.\text{right } f'; \\ & \quad g', f'' \\ & \text{ fi} \end{aligned} \quad (23)$$

The specification is very similar to *copydag*'s.

$$\begin{aligned} & \{\text{graph } g\ \gamma\ U\ V \wedge \text{ran } U = \text{dom } f\} \\ & \quad g', f' := \text{copygraph } g\ f \\ & \{\text{graph } g\ \gamma\ U\ V \star \text{graph } g'\ \gamma\ (f \bullet U) (f' \bullet V) \wedge \text{ran } V = \text{dom } f'\} \end{aligned} \quad (24)$$

The proof is extremely similar too, so it's omitted.

This problem was relatively easy to solve, once we had partial graphs. It would have been very difficult to solve had we not first analysed DAGs: instead of single let bindings we would have needed recursive let, and we would have had to deal with parallel bindings so as to describe mutually interpenetrating subgraphs.

8.4 Too subtle?

We have striven to make our specification and our proof local, concerned only with the heap that is necessarily accessed by the procedure. In order to do that we have had to use a heap predicate that doesn't constrain all that it might, we've added part of the necessary constraint in one place rather than another, and the proof of Node copying in figure 4 is perhaps dauntingly subtle.

That would be a disappointing result if this was the end of the road, but we don't think that it is. We've achieved locality, which is a significant step forward. We can glimpse ways in which other techniques – program refinement, perhaps – might make the subtlety transparent. Our proofs are more subtle than we would wish, but we are by no means disappointed.

9 Graphs with forwarding pointers

It's clear from the development so far that graphs are no harder to deal with than DAGs. In some ways they are easier: the heap predicates are simpler and have fewer cases, for example. Algorithms

that deal with cyclic graphs will certainly deal with DAGs, and now that we have the partial graph notion, we need not consider DAGs any further.

9.1 Graph copying with forwarding pointers

The algorithm of (23) searches an association list as it encounters each subgraph. That takes time, and the association list takes space – at least three pointers per entry, or four if you aren't careful or you do it in ML. It's much faster, and more efficient in space if graphs are ever to be copied or disposed, to store a forwarding pointer in each node as it is copied. Once again, we started with an algorithm:

$$\begin{aligned} \text{copyfwdgraph } g \hat{=} & \text{ if } g = \text{nil} \text{ then nil} \\ & \text{ elsef } g.\text{fwd} \neq \text{nil} \text{ then } g.\text{fwd} \\ & \text{ else} \\ & \quad g.\text{fwd} := \text{new}(\text{nil}, g.\text{val}, \text{nil}, \text{nil}); \\ & \quad g.\text{fwd}.\text{left} := \text{copyfwdgraph } g.\text{left}; \\ & \quad g.\text{fwd}.\text{right} := \text{copyfwdgraph } g.\text{right}; \\ & \quad g.\text{fwd} \\ & \text{ fi} \end{aligned} \quad (25)$$

Nodes in the heap now have an extra slot for a forwarding pointer. Uncopied subgraphs must have nil forwarding pointers; copied subgraphs point to their copy. To cut cycles we must copy the root of a new subgraph before copying its children (first line of the else case); necessarily the copy must initially have dummy subgraphs, replaced by copies as they are made.

Given a 'clean' graph – one with nil forwarding pointers – this algorithm will produce a clean copy and a separate 'dirty' original – one whose forwarding pointers define an isomorphism between the two graphs. (If you want the clean original back you have to traverse the input graph again with a cleanup procedure which we don't discuss here.) As before, we approach this ideal specification slowly, but this time armed with the partial graph notation.

The heap representation predicate of our graphs (*fwdgraph*) is identical to that of figure 5 except that nodes are now $g \mapsto _, \alpha, l, r$, leaving a slot for a forwarding pointer or (as we shall see) a disposal link.

The specification of (25) needs care if we are to observe separation between the different regions of the heap. There is no longer a forwarding function argument, but the same forwarding information is present, distributed on input across those parts of the graph that have already been copied – $\text{ran } U$ in the heap predicate – and on output distributed across $\text{ran } V$. Those nodes of the input graph which have not yet been copied – $(\text{ran}(V - U))$ – must have nil forwarding pointers, and are part of the input graph in the same way that the fringe list was part of the output tree in section 5.1. The specification (figure 6) is intricate, but not unreasonably so. Because forwarding pointers outside the partial-graph heap are explicitly mentioned and must be accessible to the procedure, we don't need to say that $\text{ran } U$ dangles outside that heap: it obviously does.

The proof is straightforward. The Empty case (figure 7) is immediate. In the Ptr case (figure 8) we know that $g \in \text{ran } V$, and that $U \subseteq V$ – both from the heap predicate – and so $g.\text{fwd} \neq \text{nil}$, together with the precondition, tells us that $g \in \text{ran } U$. In presenting the GNode case (figure 9) we have compiled multi-step instructions into sequences of assignment, run together some compound assignments and omitted to indicate obvious uses of the frame rule.

The specification does support our ideal specification. A clean

$$\begin{aligned}
& \{ (\text{fwdgraph } g \ \gamma U \ V \wedge ((\forall_{*x} \in \text{ran}(V-U) \cdot x \mapsto \text{nil}) \star \text{True})) \star (\forall_{*y} \in \text{ran } U \cdot y \mapsto f \ y) \wedge \text{dom } f = \text{ran } U \wedge \text{nil} \notin \text{ran } f \} \\
& \text{if } g = \text{nil} \text{ then} \\
& \quad \{ (\text{fwdgraph nil Empty } U \ U \wedge ((\forall_{*x} \in \text{ran}(U-U) \cdot x \mapsto \text{nil}) \star \text{True})) \star (\forall_{*y} \in \text{ran } U \cdot y \mapsto f \ y) \wedge \text{dom } f = \text{ran } U \wedge \text{nil} \notin \text{ran } f \} \\
& \quad \text{nil} \\
& \quad \{ (\text{fwdgraph nil Empty } U \ U \wedge ((\forall_{*x} \in \text{ran}(U-U) \cdot x \mapsto f \ x) \star \text{True})) \star \text{fwdgraph nil Empty } (f \bullet U) \ (f \bullet U) \star \} \\
& \quad \{ (\forall_{*y} \in \text{ran } U \cdot y \mapsto f \ y) \wedge \text{dom } f = \text{ran } U \wedge \text{nil} \notin \text{ran } f \}
\end{aligned}$$

Figure 7. Fast copy graph, Empty case

$$\begin{aligned}
& \{ (\text{fwdgraph } g \ \gamma U \ V \wedge ((\forall_{*x} \in \text{ran}(V-U) \cdot x \mapsto \text{nil}) \star \text{True})) \star (\forall_{*y} \in \text{ran } U \cdot y \mapsto f \ y) \wedge \text{dom } f = \text{ran } U \wedge \text{nil} \notin \text{ran } f \wedge g \neq \text{nil} \} \\
& \text{elsf } g.\text{fwd} \neq \text{nil} \text{ then} \\
& \quad \left\{ \begin{array}{l} (\text{fwdgraph } g \ \gamma U \ V \wedge ((\forall_{*x} \in \text{ran}(V-U) \cdot x \mapsto \text{nil}) \star \text{True})) \star (\forall_{*y} \in \text{ran } U \cdot y \mapsto f \ y) \wedge \\ \text{dom } f = \text{ran } U \wedge \text{nil} \notin \text{ran } f \wedge g \in \text{ran } U \end{array} \right\} \dots \\
& \quad \left\{ \begin{array}{l} (\text{fwdgraph } g \ (\text{Ptr } x) \ U \ U \wedge ((\forall_{*x} \in \text{ran}(U-U) \cdot x \mapsto \text{nil}) \star \text{True})) \star (\forall_{*y} \in \text{ran } U \cdot y \mapsto f \ y) \wedge \\ \text{dom } f = \text{ran } U \wedge \text{nil} \notin \text{ran } f \wedge g \in \text{ran } U \end{array} \right\} \\
& \quad \left\{ (\text{fwdgraph } g \ (\text{Ptr } x) \ U \ U \wedge ((\forall_{*x} \in \text{ran}(U-U) \cdot x \mapsto f \ x) \star \text{True})) \star \text{fwdgraph } (f \ g) \ (\text{Ptr } x) \ (f \bullet U) \ (f \bullet U) \star \right\} \\
& \quad \{ (\forall_{*y} \in \text{ran } U \cdot y \mapsto f \ y) \wedge \text{dom } f = \text{ran } U \wedge \text{nil} \notin \text{ran } f \} \\
& \quad g.\text{fwd} \\
& \quad \left\{ \begin{array}{l} (\text{fwdgraph } g \ (\text{Ptr } x) \ U \ U \wedge ((\forall_{*x} \in \text{ran}(U-U) \cdot x \mapsto f \ x) \star \text{True})) \star \text{fwdgraph } \text{res} \ (\text{Ptr } x) \ (f \bullet U) \ (f \bullet U) \star \\ (\forall_{*y} \in \text{ran } U \cdot y \mapsto f \ y) \wedge \text{dom } f = \text{ran } U \wedge \text{nil} \notin \text{ran } f \end{array} \right\}
\end{aligned}$$

Figure 8. Fast copy graph, Ptr case

$$\begin{aligned}
& \left\{ \begin{array}{l} (\text{fwdgraph } g \ (x : \text{GNode } \alpha \ \lambda \ \rho) \ U \ V \wedge ((\forall_{*x} \in \text{ran}(V-U) \cdot x \mapsto \text{nil}) \star \text{True})) \star (\forall_{*y} \in \text{ran } U \cdot y \mapsto f \ y) \wedge \\ \text{dom } f = \text{ran } U \wedge \text{nil} \notin \text{ran } f \wedge g \neq \text{nil} \end{array} \right\} \dots \\
& \left\{ \begin{array}{l} ((g \mapsto \text{nil}, \alpha, l, r \star \text{fwdgraph } l \ \lambda \ (U \oplus (x : g)) \ U' \star \text{fwdgraph } r \ \rho \ U' \ V) \wedge ((\forall_{*x} \in \text{ran}(V-U) \cdot x \mapsto \text{nil}) \star \text{True})) \star \\ (\forall_{*y} \in \text{ran } U \cdot y \mapsto f \ y) \wedge \text{dom } f = \text{ran } U \wedge \text{nil} \notin \text{ran } f \end{array} \right\} \\
& g' := \text{new}(\text{nil}, g.\text{val}, \text{nil}, \text{nil}); \\
& \left\{ \begin{array}{l} ((g \mapsto \text{nil}, \alpha, l, r \star \text{fwdgraph } l \ \lambda \ (U \oplus (x : g)) \ U' \star \text{fwdgraph } r \ \rho \ U' \ V) \wedge ((\forall_{*x} \in \text{ran}(V-U) \cdot x \mapsto \text{nil}) \star \text{True})) \star \\ g' \mapsto \text{nil}, \alpha, \text{nil}, \text{nil} \star (\forall_{*y} \in \text{ran } U \cdot y \mapsto f \ y) \wedge \text{dom } f = \text{ran } U \wedge \text{nil} \notin \text{ran } f \end{array} \right\} \\
& g.\text{fwd} := g'; \\
& \left\{ \begin{array}{l} ((g \mapsto g', \alpha, l, r \star \text{fwdgraph } l \ \lambda \ (U \oplus (x : g)) \ U' \star \text{fwdgraph } r \ \rho \ U' \ V) \wedge (g \mapsto g' \star (\forall_{*x} \in (\text{ran}(V-U) - \{g\}) \cdot x \mapsto \text{nil}) \star \text{True})) \star \\ (\forall_{*y} \in \text{ran } U \cdot y \mapsto f \ y) \star g' \mapsto \text{nil}, \alpha, \text{nil}, \text{nil} \wedge \text{dom } f = \text{ran } U \wedge \text{nil} \notin \text{ran } f \end{array} \right\} \\
& l' := \text{copyfwdgraph } g.\text{left}; \\
& \left\{ \begin{array}{l} ((g \mapsto g', \alpha, l, r \star \text{fwdgraph } l \ \lambda \ (U \oplus (x : g)) \ U' \star \text{fwdgraph } r \ \rho \ U' \ V) \wedge \\ ((\forall_{*x} \in \text{ran}(U' - U) \cdot x \mapsto f' \ x) \star (\forall_{*x} \in \text{ran}(V - U') \cdot x \mapsto \text{nil}) \star \text{True})) \star (\forall_{*y} \in \text{ran } U \cdot y \mapsto f \ y) \star g' \mapsto \text{nil}, \alpha, \text{nil}, \text{nil} \star \\ (\text{fwdgraph } l' \ \lambda \ ((f \oplus (g : g')) \bullet (U \oplus (x : g))) \ ((f \oplus f') \bullet U')) \wedge \text{dom } f = \text{ran } U \wedge \text{nil} \notin \text{ran } f \wedge \text{dom } f' = \text{ran } U' \wedge \text{nil} \notin \text{ran } f' \end{array} \right\} \\
& g.\text{fwd}.\text{left} := l'; \\
& \left\{ \begin{array}{l} ((g \mapsto g', \alpha, l, r \star \text{fwdgraph } l \ \lambda \ (U \oplus (x : g)) \ U' \star \text{fwdgraph } r \ \rho \ U' \ V) \wedge \\ ((\forall_{*x} \in \text{ran}(U' - U) \cdot x \mapsto f' \ x) \star (\forall_{*x} \in \text{ran}(V - U') \cdot x \mapsto \text{nil}) \star \text{True})) \star (\forall_{*y} \in \text{ran } U \cdot y \mapsto f \ y) \star g' \mapsto \text{nil}, \alpha, l', \text{nil} \star \\ (\text{fwdgraph } l' \ \lambda \ ((f \oplus (g : g')) \bullet (U \oplus (x : g))) \ ((f \oplus f') \bullet U')) \wedge \text{dom}(f \oplus f') = \text{ran } U' \wedge \text{nil} \notin \text{ran}(f \oplus f') \end{array} \right\} \\
& r' := \text{copyfwdgraph } g.\text{right}; \\
& \left\{ \begin{array}{l} ((g \mapsto g', \alpha, l, r \star \text{fwdgraph } l \ \lambda \ (U \oplus (x : g)) \ U' \star \text{fwdgraph } r \ \rho \ U' \ V) \wedge \\ ((\forall_{*x} \in \text{ran}(U' - U) \cdot x \mapsto f' \ x) \star (\forall_{*x} \in \text{ran}(V - U') \cdot x \mapsto f'' \ x) \star \text{True})) \star (\forall_{*y} \in \text{ran } U \cdot y \mapsto f \ y) \star g' \mapsto \text{nil}, \alpha, l', \text{nil} \star \\ (\text{fwdgraph } l' \ \lambda \ ((f \oplus (g : g')) \bullet U \oplus (x : g)) \ ((f \oplus f') \bullet U')) \star (\text{fwdgraph } r' \ \rho \ ((f \oplus f') \bullet U') \ ((f \oplus f' \oplus f'') \bullet V)) \wedge \\ \text{dom}(f \oplus f' \oplus f'') = \text{ran } V \wedge \text{nil} \notin \text{ran}(f \oplus f' \oplus f'') \end{array} \right\} \\
& g.\text{fwd}.\text{right} := r'; \\
& \left\{ \begin{array}{l} ((g \mapsto g', \alpha, l, r \star \text{fwdgraph } l \ \lambda \ (U \oplus (x : g)) \ U' \star \text{fwdgraph } r \ \rho \ U' \ V) \wedge \\ ((\forall_{*x} \in \text{ran}(V-U) \cdot x \mapsto (f \oplus f'') \ x) \star \text{True})) \star (\forall_{*y} \in \text{ran } U \cdot y \mapsto f \ y) \star g' \mapsto \text{nil}, \alpha, l', r' \star \\ (\text{fwdgraph } l' \ \lambda \ ((f \oplus (g : g')) \bullet U \oplus (x : g)) \ ((f \oplus f') \bullet U')) \star (\text{fwdgraph } r' \ \rho \ ((f \oplus f') \bullet U') \ ((f \oplus f' \oplus f'') \bullet V)) \wedge \\ \text{dom}(f \oplus f' \oplus f'') = \text{ran } V \wedge \text{nil} \notin \text{ran}(f \oplus f' \oplus f'') \end{array} \right\} \dots \\
& \left\{ \begin{array}{l} (\text{fwdgraph } g \ (\text{GNode } \alpha \ \lambda \ \rho) \ U \ V \wedge ((\forall_{*x} \in \text{ran}(V-U) \cdot x \mapsto (f' \oplus f'') \ x) \star \text{True})) \star \\ (\forall_{*y} \in \text{ran } U \cdot y \mapsto f \ y) \star \text{fwdgraph } (f'' \ g) \ (\text{GNode } \alpha \ \lambda \ \rho) \ (f \bullet U) \ (f'' \bullet V) \wedge \text{dom}(f \oplus f' \oplus f'') = \text{ran } V \wedge \text{nil} \notin \text{ran}(f \oplus f' \oplus f'') \end{array} \right\} \\
& g.\text{fwd} \\
& \left\{ \exists f''' \cdot \left((\text{fwdgraph } g \ (\text{GNode } \alpha \ \lambda \ \rho) \ U \ V \wedge ((\forall_{*x} \in \text{ran}(V-U) \cdot x \mapsto f''' \ x) \star \text{True})) \star (\forall_{*y} \in \text{ran } U \cdot y \mapsto f \ y) \star \right) \right\} \\
& \left\{ \begin{array}{l} (\text{fwdgraph } \text{res} \ (\text{GNode } \alpha \ \lambda \ \rho) \ (f \bullet U) \ (f''' \bullet V) \wedge \text{dom}(f \oplus f''') = \text{ran } V \wedge \text{nil} \notin \text{ran}(f \oplus f''')) \\ (\forall_{*y} \in \text{ran } U \cdot \exists z \cdot (y \mapsto z \wedge z \neq \text{nil})) \end{array} \right\}
\end{aligned}$$

Figure 9. Fast copy graph, GNode case

$$\begin{aligned}
& \{ (\text{fwdgraph } g \ \gamma U \ V \wedge ((\forall_{*x} \in \text{ran}(V-U) \cdot x \mapsto \text{nil}) \star \text{True})) \star (\forall_{*y} \in \text{ran } U \cdot \exists z \cdot (y \mapsto z \wedge z \neq \text{nil})) \wedge g \neq \text{nil} \rightarrow c \neq \text{nil} \} \\
& \quad c' := \text{linknodes } g \ c \\
& \left\{ \begin{array}{l} (((c' = c \wedge \text{emp}) \vee (\text{nseg } c' \ g \ _ \text{nseg } g \ c \ (g))) \wedge ((\forall_{*x} \in \text{ran}(V-U) \cdot \exists z' \cdot (x \mapsto z' \wedge z' \neq \text{nil})) \star \text{True})) \star \\ (\forall_{*y} \in \text{ran } U \cdot \exists z \cdot (y \mapsto z \wedge z \neq \text{nil})) \end{array} \right\}
\end{aligned}$$

Figure 10. Specification of linknodes

graph is

$$\text{fwdgraph } g \ \gamma \text{null } V \wedge ((\forall_* x \in \text{ran } V \cdot x \mapsto \text{nil}) \star \text{True})$$

and a dirty graph is

$$\text{fwdgraph } g \ \gamma \text{null } V \wedge ((\forall_* x \in \text{ran } V \cdot \exists y \cdot (x \mapsto y \wedge y \neq \text{nil})) \star \text{True})$$

The overall specification is

$$\{\text{clean } g \ \gamma\} \ g' := \text{copyfwdgraph } g \ \{\text{dirty } g \ \gamma \star \text{clean } g' \ \gamma\} \quad (26)$$

9.2 Graph disposal with forwarding pointers

Graph disposal using forwarding pointers takes more ingenuity than graph copying, because a program can't access already-disposed structures: that is, we can't dispose of the graph in a recursive scan without using an auxiliary list as in *disposegraph* above. The trick is to link together the nodes of the graph using the forwarding pointer slots.

$$\begin{aligned} \text{linknodes } g \ c \hat{=} & \text{ if } g = \text{nil} \vee \vee g.\text{fwd} \neq \text{nil} \text{ then } c \\ & \text{ else} \\ & \quad g.\text{fwd} := c; \\ & \quad c' := \text{linknodes } g.\text{left } g; \\ & \quad c'' := \text{linknodes } g.\text{right } g; \\ & \quad c'' \\ & \text{ fi} \end{aligned} \quad (27)$$

The second trick is to make a list that doesn't end in nil, because nil is used to distinguish unvisited nodes. Rejecting the temptation to invent another kind of nil – no true programmer would succumb! – we construct a sequence in which the last element points to itself. Then *disposenodes* (*linknodes* *g* *g*) does both tricks.

$$\begin{aligned} \text{disposenodes } g \hat{=} & \text{ if } g \neq \text{nil} \text{ then} \\ & \quad \text{do } \text{oldg} := g; \ g := g.\text{fwd}; \ \text{dispose } \text{oldg} \ 4 \\ & \quad \text{until } g = \text{oldg} \\ & \text{ fi} \end{aligned} \quad (28)$$

In effect (27) 'linearises' the graph, converting a structure in which there can be many pointers to a node into a list in which there is only one. But it doesn't really linearise it: all the subgraph pointers are still there, but we simply ignore them; and the loop at the end of the 'list' means that it isn't a list at all. Truly, separation logic provers embrace the dangling pointer!

The specification of *linknodes* (figure 10) uses the same sort of approach as figure 6 and appeals to a heap predicate *nseg* which is analogous to *lseg*, but which represents a list of node addresses.

$$\begin{aligned} \text{nseg } c \ c \ \langle \rangle & \hat{=} \text{emp} \\ \text{nseg } g \ c \ (\langle g \rangle ++ gs) & \hat{=} \exists g'. (g \mapsto g', _, _, _ \star \text{nseg } g' \ c \ gs) \end{aligned} \quad (29)$$

Proof is omitted.

10 Summary

We have attempted to extend the reach of separation logic by considering difficult examples, and to a considerable extent we have succeeded. We have found an idiomatic form, $A \wedge (B \star \text{True})$, for inclusion sharing. We have discovered 'partial graphs' in the grammar of (22). Our proofs have punched a hole in a significant obstacle: until now these problems hadn't been dealt with by local reasoning.

This is by no means the last word, or even our last word, on the subject. One of our aims is to find ways in which pointer-mutating programs, their data structure and their proofs come apart and compose along the same seams, so that we can begin to develop a structured programming for mutated pointer structures with sharing. In the fringe linking example we feel that we have achieved that aim. In the case of algorithms even as simple as those which copy and dispose DAGs and graphs, we can see that there is more work to do. The specifications require too much information about the environment *U* and its relation to *V* – by contrast the *fringelink* argument *c* can be any pointer at all.

We are particularly excited about the potential of partial graphs, both as a syntactic specification mechanism and as a semantics for graphs. We would like to know if the idea which we developed to solve a particular problem might have deeper more useful mathematical properties.

We hope that further work will result in more than just a number of tidier proofs. It may lead to a greater understanding of graphs and how to reason about them in general. It surely ought to help us to extend the reach of separation logic to larger examples.

Acknowledgements

This work springs from endless discussion within the London separation logic community between ourselves, Josh Berdine, Ivana Mijajlovic and Noah Torp-Smith. We are particularly grateful to Torp-Smith for inspiring us, by his proof of the Cheney algorithm, to work again on graphs and for loan of the iterated conjunction operator. We are grateful for advice from John Reynolds and Hongsoek Yang. Christian Huyck pointed out to one of us how simple and fast graph copying is if you use forwarding pointers. Tony Hoare and Philippa Gardner provided useful advice on presentation of specifications and proofs. An anonymous reviewer spotted a bug in the slow *copydag* and *disposedag* algorithms shown in (14), which forced us to confront explicitly the peculiarities of our heap predicates.

11 References

- [1] R. Bornat. Proving pointer programs in Hoare logic. In R. C. Backhouse and J. N. Oliveira, editors, *Mathematics of Program Construction, 5th International Conference*, LNCS, pages 102–126. Springer, 2000.
- [2] L. Cardelli, P. Gardner, and G. Ghelli. Querying trees with pointers. Unpublished notes, 2003.
- [3] C. J. Cheney. A nonrecursive list compacting algorithm. *Comm. ACM*, 13(11):677–678, 1970.
- [4] S. Ishtiaq and P. O'Hearn. BI as an assertion language for mutable data structures. In *28th POPL*, pages 14–26, London, January 2001.
- [5] P. O'Hearn, J. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In L. Fribourg, editor, *CSL 2001*, pages 1–19. Springer-Verlag, 2001. LNCS 2142.
- [6] P. W. O'Hearn and D. J. Pym. The logic of bunched implications. *Bulletin of Symbolic Logic*, 5(2):215–244, June 1999.
- [7] D. Pym. *The Semantics and Proof Theory of the Logic of Bunched Implications*, volume 26 of *Applied Logic Series*. Kluwer Academic Publishers, 2002.

- [8] J. Reynolds. Separation logic: a logic for shared mutable data structures. Invited Paper, LICS'02, 2002.
- [9] J. C. Reynolds. Intuitionistic reasoning about shared mutable data structure. In J. Davies, B. Roscoe, and J. Woodcock, editors, *Millennial Perspectives in Computer Science*, pages 303–321. Palgrave, 2000.
- [10] H. Schorr and W. Waite. An efficient machine independent procedure for garbage collection in various list structures. *Comm. ACM*, 10(8):501–506, 1967.
- [11] N. Torp-Smith. *Proving a garbage collector using local reasoning*. Master's thesis, IT University of Copenhagen, 2003.
- [12] H. Yang. An example of local reasoning in BI pointer logic: the Schorr-Waite graph marking algorithm. Proceedings of the SPACE Workshop, 2001.
- [13] H. Yang. *Local Reasoning for Stateful Programs*. Ph.D. thesis, University of Illinois, Urbana-Champaign, Illinois, USA, 2001.