

Small Specifications for Tree Update

Philippa Gardner and Mark Wheelhouse

Imperial College London, {pg, mjw03}@doc.ic.ac.uk

Abstract. O’Hearn, Reynolds and Yang introduced Separation Logic to provide modular reasoning about simple, mutable data structures in memory. They were able to construct small specifications of programs, by reasoning about the local parts of memory accessed by programs. Gardner, Calcagno and Zarfaty generalised this work, introducing Context Logic to reason about more complex data structures. In particular, they developed a formal, compositional specification of the Document Object Model, a W3C XML update library. Whilst keeping to the spirit of local reasoning, they were not able to retain small specifications. We introduce Segment Logic, which provides a more fine-grained analysis of the tree structure and yields small specifications. As well as being aesthetically pleasing, small specifications are important for reasoning about concurrent tree update.

1 Introduction

Separation Logic [13], introduced by O’Hearn, Reynolds and Yang, provides modular reasoning about mutable data structures in memory. The idea is to reason about the small, local parts of memory (the footprint) that are accessed by a program. The resulting modular reasoning has been used to notable success for verifying memory safety properties of large C-programs [1], and for reasoning about concurrent imperative programs [14]. Calcagno, Gardner and Zarfaty generalised this work to more complex data structures, such as those found on the Web, by introducing Context Logic for reasoning about arbitrary structured data update [3]. Their original work applied Context Logic reasoning to a simple tree update language. With Smith and Zarfaty, Gardner and Wheelhouse have since applied Context Logic reasoning to the W3C Document Object Model (DOM) [6], a library for in-place XML update [18].

Our goal is to design and formally specify a concurrent XML update language. Such a language will enable web applications to make the most of the dynamic nature of XML. For example, with Wikipedia, users currently copy articles on to their browsers, before updating and returning them to Wikipedia to be integrated with the main site. They cannot view Wikipedia (or a scientific database or information in the Cloud) as a shared XML memory store that can be concurrently updated by many clients, because methods for safely performing such operations are poorly understood. We almost have the technology to develop a safe, formally specified language for concurrent XML update, drawing on our experience with sequential DOM [6] and O’Hearn’s Concurrent Separation Logic [14]. However, we are missing one ingredient.

In our DOM work, we were not able to provide small specifications for all our DOM programs. In particular, our reasoning for the basic move commands,




Fig. 1. Splitting up the Working Tree using Multi-holed Contexts.

such as DOM's `appendChild`, used axioms which required a substantial over-approximation of the footprint. Whilst this over-approximation was acceptable for reasoning about sequential programs, it is a serious limitation when reasoning about concurrent programs. In this paper, we solve this limitation, by introducing Segment Logic to provide a more fine-grained analysis of structured data update in general, and tree update in particular. We provide small axioms for all the basic commands of a simple tree update language; it is straightforward to extend our ideas to DOM [6]. Although this paper focuses on a sequential tree update language, we believe it provides the technology necessary for our future work on reasoning about concurrent tree update.

To motivate Segment Logic, consider the DOM command `appendChild(n, m)` which moves the tree with top node identified by m to be the last child of the tree identified by n . Fig. 1 indicates how the working tree splits in the two cases where `appendChild(n, m)` does not fault: it succeeds when n and m are in different parts of the tree and when m is under n ; it faults when m is above n . The axiom for `appendChild(n, m)` using multi-holed Context Logic [2] is ¹:

$$\begin{aligned} & \{(C \circ_{\alpha} n[c_1]) \circ_{\beta} m[\text{tree}(c_2)]\} \\ & \quad \text{appendChild}(n, m) \\ & \{(C \circ_{\alpha} n[c_1 \otimes m[\text{tree}(c_2)]]\} \circ_{\beta} \emptyset \end{aligned}$$

The precondition specifies that the working tree can be split into a subtree with top node m , and a tree context with hole variable β (y in Fig. 1) satisfying the separating application formula $C \circ_{\alpha} n[c_1]$. This formula states that the context can be further split into a subcontext with top node n and an unspecified context with hole α (x in Fig. 1) given by context variable C . The postcondition states that the tree at m is moved to be the last child of n and is replaced by the empty tree. The surrounding context, denoted by variable C , remains the same.

The problem with this `appendChild` axiom is that it is not small. The precondition is not the intuitive footprint. The only part of the tree that `appendChild(n, m)` requires is the tree at m which is being moved, and the tree or context with top node n (actually node n is enough) whose children are being extended by m . However, our precondition does not just use m and n . It also requires the surrounding context denoted by C . It is possible to put additional constraints on C

¹ In [6], the axiom for `appendChild` is given using single-holed Context Logic. The multi-holed Context Logic axiom is simpler, but still not suitable for concurrent reasoning.




Fig. 2. Splitting up the Working Tree using Tree Segments.

to insist that the context is minimal. But this is not the point. We need a finer way of analysing the tree in order to capture the intuitive footprint of the command.

Instead of basing our reasoning on multi-holed tree contexts and application, we base our reasoning on *tree segments*. With multi-holed contexts, the working tree is split into a context and subtrees which have lost the information about where they originated from; the application function determines which holes get filled. With segments, the working tree is split into tree segments which still ‘know’ how to join back together again. As well as unique hole labels, tree segments have unique hole addresses which determine which holes the segments fill. For example, consider Fig. 2. In both cases, the working tree is split into a bunch of tree segments. The hole labels (in the holes) and the hole addresses (on the arrows) determine how the tree segments join back up to form the original tree. Notice that hole labels and addresses have brackets around them, denoting that they are bound. In the syntax, we will use a hiding operator (x) , analogous to the restriction operator of Milner’s π -calculus [12].

Moreover, consider the right-hand equalities of Fig. 2. In both cases, the tree segment with top node identified by n has been split into just the node n at the same address and fresh hole label z , plus another tree segment at address z which contains the children of node n . We shall see that the node n and the tree with top node m are all that is required to provide the small axiom for `appendChild`. Fig. 2 thus indicates how we can uniformly separate the minimal data required in order to reason about `appendChild`. It is possible to take this separation to the extreme, by cutting up the tree structure into a collection of nodes, with the hole labels and addresses showing how the nodes are joined together (a spaghetti of wires analogous to a heap representation). However, this is not how we use the hole information. We only cut up the tree in a minimal way in order to provide the right segment about which to reason.

We introduce Segment Logic for reasoning about our tree segments. It is like Context Logic in that it reasons directly about high-level trees. It is like Separation Logic in that it uses a commutative separating conjunction $*$, rather than the non-commutative separating application of Context Logic. Using Segment Logic, the small axiom for `appendChild`(n, m) is:

$$\begin{aligned} & \{ \alpha \leftarrow n[\gamma] * \beta \leftarrow m[\text{tree}(c)] \} \\ & \text{appendChild}(n, m) \\ & \{ \alpha \leftarrow n[\gamma \otimes m[\text{tree}(c)]] * \beta \leftarrow \emptyset_T \} \end{aligned}$$

The precondition specifies two tree segments: a node n at address variable α (x in Fig. 2) and a complete tree whose top node is m at address β (y in Fig 2). The postcondition states that the tree at m moves to be the last child of n and is replaced by the empty tree. The axiom is small, with the precondition capturing the intuitive footprint of `appendChild(n, m)`. We can extend the axiom to larger tree segments using the normal separation frame rule, a rule for the hiding quantification, and the rule for logical consequence. In fact, instead of using the hiding quantifier as primitive, we use the basic revelation connective (and a revelation frame rule), the revelation magic wand (the revelation right adjoint), and the fresh label quantification (and a fresh variable elimination rule), inspired by the work of Gabbay and Pitts [5], and Cardelli and Gordon [4]. Interestingly, we shall see that these more primitive constructs are important for describing the weakest preconditions.

2 Tree Update Language

We study a simple, but expressive, high-level tree update language for manipulating finite, ordered, unranked trees, with unique node identifiers for specifying the locations of updates as in DOM. Our tree structures are left intentionally simple. It is straightforward to incorporate (and reason about) additional data such as text data and attributes (see [6]). To simplify our exposition, we work with multi-holed tree contexts [2]. Throughout this paper we use countably infinite and disjoint sets $I = \{m, n, \dots\}$ for location names and $X = \{x, y, z, \dots\}$ for hole identifiers.

Definition 1 (Tree Contexts). Multi-holed tree contexts $c \in C_{I, X}$ are defined by:

$$\begin{aligned} \text{tree context } c ::= & \emptyset_C \text{ empty tree context} \\ & x \text{ hole identifier } x \text{ used as a hole label} \\ & n[c] \text{ tree context with top node } n \\ & c \otimes c \text{ composition of tree contexts} \end{aligned}$$

with the restriction that each hole identifier, $x \in X$, and location name, $n \in I$, occur at most once in a tree context c , and subject to an equivalence $c_1 \equiv c_2$ stating that \otimes is associative with identity \emptyset_C . The set of hole identifiers that occur in tree context c is denoted $\text{free}(c)$. A tree context with no context holes (a complete tree) is denoted t .

Definition 2 (Context Application). Context Application is defined as a set of partial functions $ap_x : C_{I, X} \times C_{I, X} \rightarrow C_{I, X}$ indexed by hole labels x :

$$ap_x(c_1, c_2) = \begin{cases} c_1[c_2/x] & \text{if } x \in \text{free}(c_1) \text{ and } \text{free}(c_1) \cap \text{free}(c_2) = \{\} \\ \text{undefined} & \text{otherwise} \end{cases}$$

We abbreviate $ap_x(c_1, c_2)$ by $c_1 \circ_x c_2$. We often omit the \emptyset_C leaves from a tree context to make it more readable, writing $n[m \otimes p]$ instead of $n[m[\emptyset_C] \otimes p[\emptyset_C]]$.

Our update language is a high-level, stateful, sequential, imperative language, based on variable assignment and update commands as in DOM. The program state is made up of two components: the working tree which contains all of the nodes we will be manipulating with our programs; and a high-level variable store containing variables for node identifiers.

Definition 3 (Variable Store). The variable store $\sigma \in \Sigma$ is a finite partial function

$$\sigma : \text{Var}_1 \rightarrow_{\text{fin}} \text{I} \cup \{\mathbf{null}\}$$

mapping location name variables $\text{Var}_1 = \{m, n, \dots\}$ to location names or **null**. We write $\sigma[n \mapsto n]$ for the variable store σ overwritten with $\sigma(n) = n$.

To specify location name values, our language uses simple expressions. Location names are specified either with location name variables or the constant **null**; we forbid direct reference to constant location names other than **null**. We also require simple Boolean expressions for conditional tests in our language.

Definition 4 (Expressions). Location name expressions $N \in \text{Exp}_1$ and Boolean expressions $B \in \text{Exp}_B$ are defined by:

$$\begin{aligned} N &::= n \mid \mathbf{null} & n &\in \text{Var}_1 \\ B &::= N = N \mid \mathbf{false} \mid B \Rightarrow B \end{aligned}$$

The valuation of an expression E in a store σ is written $\llbracket E \rrbracket \sigma$ and has the obvious semantics. The classical Boolean connectives **true**, \neg , \wedge and \vee are derivable.

DOM commands tend to update whole trees although, for example, the DOM commands `getNodeName` and `createNode` manipulate single nodes. Since our reasoning analyses tree segments, we explore a language for manipulating tree segments with primitive commands that update either a single node n or the whole subtree beneath some node n ; the commands for updating whole trees are derivable (Example 1).

Definition 5 (Tree Update Language). The commands of the tree update language consist of the node update commands $\mathbb{C}_{\text{nodeUp}}$, the tree update commands $\mathbb{C}_{\text{treeUp}}$ and the standard skip, assignment, local, sequencing, if-then-else and while-do commands:

$\mathbb{C}_{\text{nodeUp}} ::= n' := \text{getUp}(n)$	get parent of node n and record it in n'
$n' := \text{getLeft}(n)$	get previous sibling of node n
$n' := \text{getRight}(n)$	get next sibling of node n
$n' := \text{getFirst}(n)$	get first child of node n
$n' := \text{getLast}(n)$	get last child of node n
<code>insertNodeAbove</code> (n)	insert a new node above node n
<code>deleteNode</code> (n)	delete node n
<code>moveNodeAbove</code> (n, m)	move node m above node n
<code>moveNodeLeft</code> (n, m)	move node m to the left of node n
<code>moveNodeRight</code> (n, m)	move node m to the right of node n
<code>prependNode</code> (n, m)	prepend node m to children of node n
<code>appendNode</code> (n, m)	append node m to children of node n
$\mathbb{C}_{\text{treeUp}} ::= \text{deleteSubtree}(n)$	delete subtree (subforest) beneath node n
<code>moveSubLeft</code> (n, m)	move children of node m to the left of node n
<code>moveSubRight</code> (n, m)	move children of node m to the right of node n
<code>prependSub</code> (n, m)	prepend children of node m to children of node n
<code>appendSub</code> (n, m)	append children of node m to children of node n

The set of free variables of a command \mathbb{C} is denoted $\text{free}(\mathbb{C})$, and the set of variables modified by \mathbb{C} is denoted $\text{mod}(\mathbb{C})$.

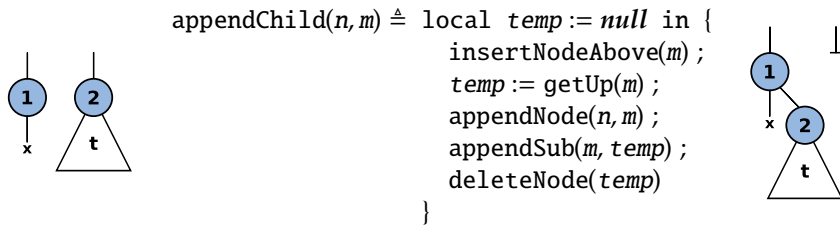
The behavior of these commands should be self-explanatory. The node update commands consist of get commands that return a neighboring node in the tree, a node insertion command that puts a fresh node into the tree above node n (the other node insertion commands are derivable), a delete command that removes a node from the tree, and node move commands that take a node out of the tree and put it in a new position. These node move commands leave the children of the moved node m as children of m 's old parent. The tree update commands work on subtrees of an identified node. They consist of a delete command that removes an entire subtree from the tree, and subtree move commands that take a subtree out of the tree and put it in a new position.

These commands are sufficient to express a wide range of tree manipulation, as illustrated by the examples below. Our command set is not minimal: for example, we could derive the `deleteSubtree` command using a combination of get commands, node deletion and recursion. However, we believe the commands chosen provide a natural and expressive tree update language. We give the operational semantics in Section 3 using tree segments, rather than trees or tree contexts, as this simplifies the description of our reasoning in Section 5.

In [8], there are also commands for inserting whole trees. This is achieved by including tree shapes (trees without identifiers) in the variable store. Here we avoid such complications by omitting commands for tree insertion and copying. The reasoning presented here extends simply to these extra commands.

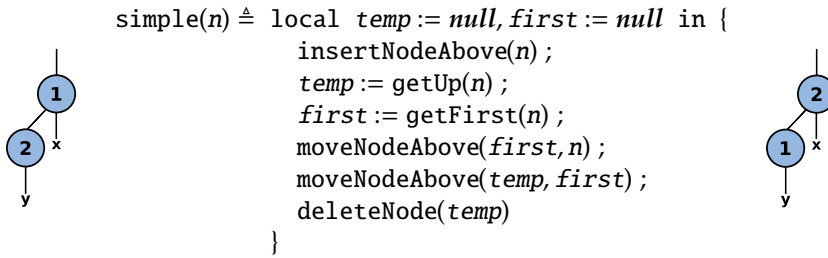
We now give example programs which will be used to illustrate our reasoning in Section 5. In all of these examples, $\sigma(n) = 1$ and $\sigma(m) = 2$.

Example 1 (Move). DOM has the command `appendChild`, whereas our language has `appendNode` and `appendSub`. We implement the standard `appendChild` as:



The diagrams illustrate the intuitive effect of the program on the working tree. For the program above not to fault, it requires the node $n = 1$ and the complete tree at node $m = 2$ (the left-hand diagram and the intuitive footprint). The complete tree at m ensures that m is not an ancestor of n . The result of the program is to move the tree at $m = 2$ to be under the node $n = 1$ (the right-hand diagram). Our reasoning captures the intuition illustrated by these diagrams.

Example 2 (Simple Swap). Our node update commands enable us to define programs that act on arbitrary segments of the tree. For example, consider the program `simple(n)` which swaps a node n with its first child:

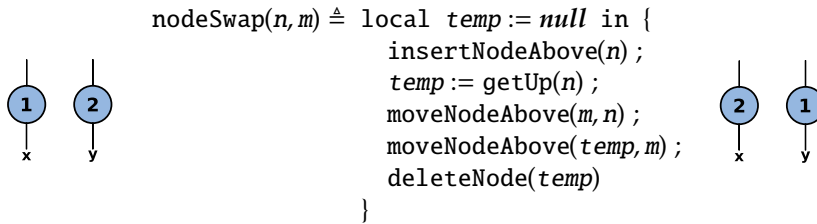


```

simple(n) ≜ local temp := null, first := null in {
  insertNodeAbove(n);
  temp := getUp(n);
  first := getFirst(n);
  moveNodeAbove(first, n);
  moveNodeAbove(temp, first);
  deleteNode(temp)
}

```

Example 3 (General Swap). The program nodeSwap(n, m) swaps the positions of arbitrary nodes n and m of a tree leaving their subtrees stationary:

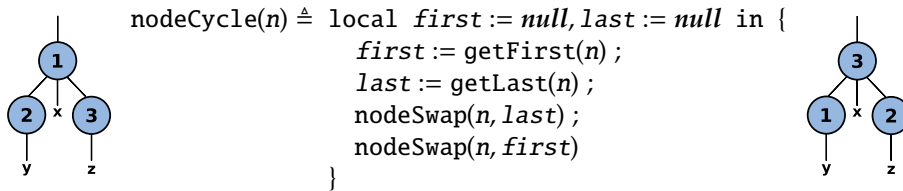


```

nodeSwap(n, m) ≜ local temp := null in {
  insertNodeAbove(n);
  temp := getUp(n);
  moveNodeAbove(m, n);
  moveNodeAbove(temp, m);
  deleteNode(temp)
}

```

Example 4 (Node Rotate). The program nodeCycle(n) takes the node n, its first and last child, and rotates these nodes with n taking the place of first child, first child taking the place of last child, and last child taking the place of n:

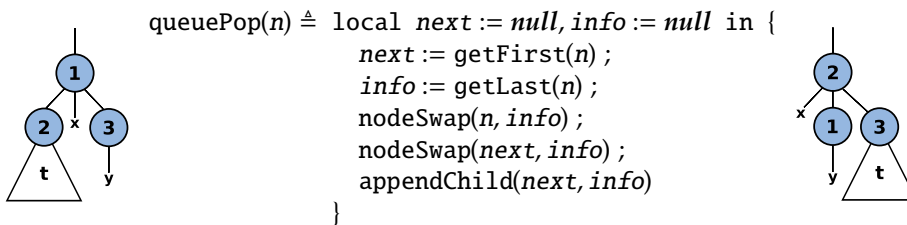


```

nodeCycle(n) ≜ local first := null, last := null in {
  first := getFirst(n);
  last := getLast(n);
  nodeSwap(n, last);
  nodeSwap(n, first)
}

```

Example 5 (Combining Move and Node Swap). Consider a simple cyclic list of pictures used, for example, to view properties on an estate agent's web page. It can be implemented as a tree structure, with the root node of the tree containing the ID of the picture currently being displayed, the picture itself being stored beneath the last of its children (under node 3 below), and the other pictures in the list being stored beneath their ID nodes as the rest of the root node's children. The program queuePop(n) cycles the pictures so that the current picture moves to the back of the list and the next picture is displayed. Notice the use of appendChild which includes the complete tree t in the footprint.



```

queuePop(n) ≜ local next := null, info := null in {
  next := getFirst(n);
  info := getLast(n);
  nodeSwap(n, info);
  nodeSwap(next, info);
  appendChild(next, info)
}

```

3 Tree Segments

We are not able to provide a small specification of the `appendChild` command (and the `appendSub` command) using tree contexts. We are able to provide small specifications for all our tree update commands using tree segments.

Definition 6 (Tree Segments). *Tree segments $s \in S_{I, X}$ are defined by the grammar:*

$$\begin{array}{ll} \text{tree segment } s ::= & \emptyset_S \text{ empty tree segment} \\ & x \leftarrow c \text{ tree context } c \text{ addressed by hole identifier } x \\ & s + s \text{ disjoint union} \\ & (x)(s) \text{ hiding, hole identifier } x \text{ bound in tree segment } s \end{array}$$

with the restriction that each hole identifier $x \in X$ occurs free at most once as a hole label and at most once as a hole address in tree segment s , and each location name, $n \in I$, occurs at most once in s . The set $\text{free}(s)$ denotes the set of free hole identifiers in s .

With tree contexts, we have the application $(1[x \otimes 3]) \circ_x 2 = 1[2 \otimes 3]$. The application \circ_x binds x , and declares that hole x is filled by the argument 2. With tree segments, we have the equivalence $(x)(z \leftarrow 1[x \otimes 3] + x \leftarrow 2) \equiv z \leftarrow 1[2 \otimes 3]$. In this case, it is the segment $x \leftarrow 2$ with address x that declares that 2 should go into the hole x , and the hiding operator (x) which binds x in the segment.

Definition 7 (Tree Segment Equivalence). *An equivalence relation \equiv over tree segments is defined by the following axioms and the natural structural rules:*

$$\begin{array}{ll} s + \emptyset_S \equiv s & s_1 + s_2 \equiv s_2 + s_1 \\ (x)(\emptyset_S) \equiv \emptyset_S & s_1 + (s_2 + s_3) \equiv (s_1 + s_2) + s_3 \\ (x)(y)(s) \equiv (y)(x)(s) & \\ (x)(s) \equiv (y)(s[y/x]) & \text{if } y \notin \text{free}(s) \\ (x)(y \leftarrow c + s) \equiv y \leftarrow c + (x)(s) & \text{if } x \neq y \text{ and } x \notin \text{free}(c) \\ (x)(y \leftarrow c_1 + x \leftarrow c_2) \equiv y \leftarrow (c_1 \circ_x c_2) & \text{if } x \in \text{free}(c_1) \end{array}$$

Most of the axioms involving hiding follow from analogous axioms for the restriction operator of the π -calculus [12]. The last hiding axiom is specific to tree segments. It enables us to pull apart and compress segments, as illustrated in **Fig. 3**. A tree segment is in its compressed form if it cannot be further compressed using this last axiom. A tree segment is well-formed if and only if its compressed form is cycle free; that is, the hole labels and hole addresses are disjoint in its compressed form. We only work with well-formed tree segments in this paper.

Fig. 3 demonstrates a graphical interpretation of segments. The left-hand side of **Fig. 3** will come as no surprise to those familiar with graphical process models: for example, Milner's work on process graphs [11]. Here, the hole identifiers describe the edges of the graph (wires). However, this is not the only use of hole identifiers. Consider the right-hand side of **Fig. 3**. Here, the hole identifier y is used to address multiple edges of a graph. More than this, consider the `appendChild` command in Example 1. The tree segment $z \leftarrow 1[x] + y \leftarrow 2[t]$ updates to $z \leftarrow 1[x \otimes 2[t]] + y \leftarrow \emptyset_C$: before update the segment $y \leftarrow 2[t]$ states that a tree is at address y ; after update $y \leftarrow \emptyset_C$ states that the empty tree is at address y .




Fig. 3. Equivalent Tree Segments: (x) and (y) denote hidden hole labels and addresses. This example illustrates that edge arity is not necessarily preserved by update. Our hole identifiers should therefore not be regarded as describing graph edges. Instead, they describe tree fragments.

Notice that our language manipulates nodes and complete trees. It does not refer to hole identifiers in any way. However, the operational semantics is greatly simplified by using either tree contexts or tree segments. We choose tree segments, as this leads to a simpler interpretation of Hoare triples in Section 5.

Definition 8 (Operational Semantics). We give the operational semantics for the basic commands of the tree update language in **Fig. 4** using an evaluation relation \rightsquigarrow relating configuration triples \mathbf{C}, σ, s , terminal states σ, s , and faults, where \mathbf{C} is a command, σ is a variable store and s is a tree segment. The set of variables of a command \mathbf{C} is denoted $\text{free}(\mathbf{C})$ and is contained within the domain of σ , denoted $\text{dom}(\sigma)$. We omit the standard cases for skip, assignment, local, sequencing, if-then-else and while-do.

Our style of local Hoare reasoning about programs requires that the commands of our language be local. A command is local if it satisfies two properties, initially introduced in [10], known as the *safety-monotonicity* property and the *frame* property. The *safety-monotonicity* property specifies that, if a command is safe (does not fault) in a given state, then it is safe in a larger state. The *frame* property specifies that, if a command is safe in a given state, then any execution on a larger state can be tracked to an execution on the smaller state. The commands of our language presented here (and the DOM commands [6]) are local. For example, consider the behavior of $n' := \text{getRight}(n)$. If the right sibling of n exists, then its identifier is stored at n' . If n is the last child of some parent node (meaning n can never obtain a right sibling via segment composition), then n' stores the value *null*. However, if the node n is not present in the tree, or n has no right sibling or parent, then the command must fault in order to be local.

4 Segment Logic

We introduce Segment Logic. First, we present the *logical environment* which contains logical variables for tree contexts, tree segments and hole identifiers. Location name variables have the standard dual role as both program variables and logical variables. They are declared in the variable store, but can be quantified like logical variables.

Definition 9 (Logical Environment). An environment $e \in \mathbf{E}$ is a set of functions

$$e : (LVar_{\mathbf{C}} \rightarrow \mathbf{C}_{I, X}) \times (LVar_{\mathbf{S}} \rightarrow \mathbf{S}_{I, X}) \times (LVar_{\mathbf{X}} \rightarrow \mathbf{X})$$

mapping tree context variables $LVar_{\mathbf{C}} = \{c, \dots\}$ to tree contexts, tree segment variables $LVar_{\mathbf{S}} = \{s, \dots\}$ to tree segments, and hole identifier variables $LVar_{\mathbf{X}} = \{\alpha, \beta, \gamma, \delta, \dots\}$ to hole identifiers. We write $e[lvar \mapsto val]$ for e overwritten with $e[lvar] = val$.

$$\begin{array}{c}
\frac{\sigma(n) = n \quad s \equiv (w, x, y, z)(s' + x \leftarrow m[y \otimes n[w] \otimes z])}{n' := \text{getUp}(n), \sigma, s \rightsquigarrow \sigma[n' \mapsto m], s} \\
\\
\frac{\sigma(n) = n \quad s \equiv (x, y, z)(s' + x \leftarrow n[y] \otimes m[z])}{n' := \text{getRight}(n), \sigma, s \rightsquigarrow \sigma[n' \mapsto m], s} \quad \frac{\sigma(n) = n \quad s \equiv (x, y, z)(s' + x \leftarrow m[z \otimes n[y]])}{n' := \text{getRight}(n), \sigma, s \rightsquigarrow \sigma[n' \mapsto \text{null}], s} \\
\\
\frac{\sigma(n) = n \quad s \equiv (x, y, z)(s' + x \leftarrow n[y \otimes m[z]])}{n' := \text{getLast}(n), \sigma, s \rightsquigarrow \sigma[n' \mapsto m], s} \quad \frac{\sigma(n) = n \quad s \equiv (x)(s' + x \leftarrow n[\otimes_C])}{n' := \text{getLast}(n), \sigma, s \rightsquigarrow \sigma[n' \mapsto \text{null}], s} \\
\\
\frac{\sigma(n) = n \quad s \equiv (x, y)(s'' + x \leftarrow n[y]) \quad m \text{ fresh id} \quad s' \equiv (x, y)(s'' + x \leftarrow m[n[y]])}{\text{insertNodeAbove}(n), \sigma, s \rightsquigarrow \sigma, s'} \quad \frac{\sigma(n) = n \quad s \equiv (x, y)(s'' + x \leftarrow n[y]) \quad s' \equiv (x, y)(s'' + x \leftarrow y)}{\text{deleteNode}(n), \sigma, s \rightsquigarrow \sigma, s'} \quad \frac{\sigma(n) = n \quad s \equiv (x)(s'' + x \leftarrow n[t]) \quad s' \equiv (x)(s'' + x \leftarrow n[\otimes_C])}{\text{deleteSubtree}(n), \sigma, s \rightsquigarrow \sigma, s'} \\
\\
\frac{\sigma(n) = n \quad \sigma(m) = m \quad s \equiv (w, x, y, z)(s'' + x \leftarrow n[z] + y \leftarrow m[w]) \quad s' \equiv (w, x, y, z)(s'' + x \leftarrow n[z \otimes m] + y \leftarrow w)}{\text{appendNode}(n, m), \sigma, s \rightsquigarrow \sigma, s'} \quad \frac{\sigma(n) = n \quad \sigma(m) = m \quad s \equiv (x, y, z)(s'' + x \leftarrow n[z] + y \leftarrow m[t]) \quad s' \equiv (x, y, z)(s'' + x \leftarrow n[z \otimes t] + y \leftarrow m)}{\text{appendSub}(n, m), \sigma, s \rightsquigarrow \sigma, s'} \quad \frac{\sigma(n) = n \quad \sigma(m) = m \quad s \equiv (w, x, y, z)(s'' + x \leftarrow n[z] + y \leftarrow m[w]) \quad s' \equiv (w, x, y, z)(s'' + x \leftarrow m[n[z]] + y \leftarrow w)}{\text{moveNodeAbove}(n, m), \sigma, s \rightsquigarrow \sigma, s'}
\end{array}$$

For get and move, only some of the cases are given; the other cases are analogous. Our commands fault when the program state does not satisfy any of the preconditions for that command.

Fig. 4. Operational Semantics for the Basic Tree Update Commands.

Segment Logic for trees consists of segment formulae and tree formulae. Just as in Separation Logic and Context Logic, segment formulae consist of classical formulae, structural formulae and specific formulae for describing the structure of data (in this case trees). For this paper, we have chosen to use tree formulae in the style of Ambient Logic [4], although we do not see a reason why P_T could not be first-order logic formulae for describing trees or even XDuce types [9]. Note that adapting this work to other data structures, such as sequences, just involves changing the tree formulae (or types) to sequence formulae (or types).

Definition 10 (Formulae). *The formulae of Segment Logic for trees consist of the segment formulae P_S and tree formulae P_T given by:*

$$\begin{array}{ll}
P_S, Q_S ::= P_S \Rightarrow P_S \mid \mathbf{false}_S & P_T, Q_T ::= P_T \Rightarrow P_T \mid \mathbf{false}_T \quad \text{Classical} \\
\mid \emptyset_S \mid P_S * P_S \mid P_S \multimap P_S \mid \alpha \textcircled{R} P_S \mid \alpha \textcircled{-R} P_S & \quad \text{Structural} \\
\mid \exists \text{var}. P_S \mid \exists \text{lvar}. P_S \mid \forall \alpha. P_S & \mid \exists \text{var}. P_T \mid \exists \text{lvar}. P_T \quad \text{Quantifiers} \\
\mid \alpha \leftarrow P_T & \mid \emptyset_T \mid \alpha \mid n[P_T] \mid P_T \otimes P_T \quad \text{Specific} \\
\mid s \mid B & \mid c \mid B \mid @_T \alpha \quad \text{Expression}
\end{array}$$

Let $\text{free}(P_S)$ and $\text{free}(P_T)$ denote the appropriate sets of free variables: α is free in $\alpha \textcircled{R} P_S$, $\alpha \textcircled{-R} P_S$, $\alpha \leftarrow P_T$ and $@_T \alpha$, bound in $\forall \alpha. P_S$. var is a location name variable and lvar is a logical variable. The binding precedence, strongest first, is: \otimes , \leftarrow , $*$, \textcircled{R} , \multimap , $\textcircled{-R}$, \Rightarrow .

The separating connective $*$, its unit \emptyset_S and its right adjoint (the separating magic wand) \multimap , are structural formulae which are known from the Separation Logic literature: formula $P_S * Q_S$ describes a segment that can be separated into a segment satisfying P_S and a disjoint segment satisfying Q_S ; formula \emptyset_S describes the empty segment; and formula $P_S \multimap Q_S$ describes a segment that, whenever it is joined to a segment satisfying P_S , results in a segment satisfying Q_S .

Before explaining the other structural formulae, we explain the specific segment formulae: the formula $\alpha \leftarrow P_T$ describes a tree segment with hole address given by the value of variable α and tree context satisfying tree formula P_T . The tree formulae follow the style of reasoning given by Ambient Logic. For




Fig. 5. Connection and Disconnection of Tree Segments.

the specific tree formulae, we have \emptyset_T specifying the empty tree context, α specifying a hole label given by the value of variable α , $n[P_T]$ specifying a tree context with top node denoted by node variable n and subtree satisfying P_T , and the composition formula $P_T \otimes Q_T$ describing a tree context which can be split into one context satisfying P_T and the other disjoint context satisfying Q_T . The tree expression formulae include the formula $@_T \alpha$ which describes a tree that contains α free; the analogous formula for tree segments is derivable.

The other structural connectives are the revelation connective, \mathbb{R} , and its right adjoint, the revelation magic wand $\neg\mathbb{R}$. As far as we are aware, these connectives have not been used in the local reasoning setting before. Together with the freshness quantifier $\mathcal{V}\alpha$, they have been used in the Ambient Logic [4], following the work of Pitts and Gabbay [5]. The freshness quantifier enables us to pick a completely new hole identifier. The formula $\alpha\mathbb{R}P_S$ describes a segment with a top-level hiding binder given by the value of α such that, after the hiding is removed, the remaining segment satisfies P_S . Consider the tree segment $z\leftarrow 1[2 \otimes 3] \equiv (x)(z\leftarrow 1[x \otimes 3] + x\leftarrow 2)$. It satisfies the formula $\alpha\mathbb{R}(\beta\leftarrow 1[\alpha \otimes \mathbf{true}_T] + \alpha\leftarrow 2)$, when $\alpha = x, \beta = z$. The revelation connective $\alpha\mathbb{R}$ strips off the hiding binder, disconnecting the tree into the fragments $z\leftarrow 1[x \otimes 3] + x\leftarrow 2$ as illustrated in Fig. 5. By contrast, the formula $\alpha\neg\mathbb{R}P_S$ describes a segment which satisfies P_S if it is extended with a hiding binder over the hole identifier stored in variable α . For example, the tree segment $z\leftarrow 1[x \otimes 3] + x\leftarrow 2$ satisfies the formula $\alpha\neg\mathbb{R}(\beta\leftarrow 1[2 \otimes 3])$ when $\alpha = x$. The revelation magic wand $\alpha\neg\mathbb{R}$ adds the binder (x) to the segment to obtain the tree segment $(x)(z\leftarrow 1[x \otimes 3] + x\leftarrow 2) \equiv z\leftarrow 1[2 \otimes 3]$, thus connecting the fragmented tree into the whole tree as illustrated in Fig. 5. Analogous to the separating magic wand, we shall see that the revelation magic wand is important for giving the weakest preconditions of commands.

Definition 11 (Satisfaction Relation). Given a logical environment e and a variable store σ , the semantics of Segment Logic is given in Fig. 6 by two satisfaction relations $e, \sigma, s \models_S P_S$ and $e, \sigma, c \models_T P_T$ defined on tree segments and tree contexts respectively.

Definition 12 (Derived Formulae). The standard classical logic connectives are derived from **false** and \Rightarrow as usual, and the following useful formulae are defined:

$$\begin{aligned} \mathit{tree}(P_T) &\triangleq P_T \wedge \neg \exists \alpha. @_T \alpha & @_S \alpha &\triangleq \mathbf{H}\beta. (\mathbf{true}_S * (\alpha \leftarrow \beta \vee \beta \leftarrow \alpha)) \\ n &\triangleq n[\emptyset_T] & \diamond P_S &\triangleq \mathbf{true}_S * P_S \\ \circ[P_T] &\triangleq \exists m. m[P_T] & \mathbf{H}\alpha. P_S &\triangleq \mathcal{V}\alpha. \alpha\mathbb{R}P_S \end{aligned}$$

Formula $\mathit{tree}(P_T)$ describes a complete tree satisfying P_T . Formula n describes a leaf node identified by n . Formula $\circ[P_T]$ allows us to drop the identifier of a node. Formula $@_S \alpha$ describes a tree segment that contains α free. Formula $\diamond P_S$ allows us to express that somewhere in the tree segment there is a segment satisfying P_S . Finally, formula $\mathbf{H}\alpha. P_S$ provides the standard hiding quantification [4] allowing us to quantify over hidden (restricted) labels.

$e, \sigma, s \models_S \emptyset_S$	$\Leftrightarrow s \equiv \emptyset_S$	$e, \sigma, c \models_{\top} \emptyset_{\top}$	$\Leftrightarrow c \equiv \emptyset_C$
$e, \sigma, s \models_S P_S * Q_S$	$\Leftrightarrow \exists s_1, s_2. s \equiv s_1 + s_2 \wedge e, \sigma, s_1 \models_S P_S \wedge e, \sigma, s_2 \models_S Q_S$	$e, \sigma, c \models_{\top} \alpha$	$\Leftrightarrow c \equiv e(\alpha)$
$e, \sigma, s \models_S P_S \multimap Q_S$	$\Leftrightarrow \forall s'. e, \sigma, s' \models_S P_S \wedge (s + s') \Downarrow \Rightarrow e, \sigma, s + s' \models_S Q_S$	$e, \sigma, c \models_{\top} n[P_{\top}]$	$\Leftrightarrow \exists c_1. c \equiv \sigma(n)[c_1]$
$e, \sigma, s \models_S \alpha @ P_S$	$\Leftrightarrow \exists x, s'. e(\alpha) = x \wedge s \equiv (x)(s') \wedge e, \sigma, s' \models_S P_S$	$e, \sigma, c \models_{\top} P_{\top} \otimes Q_{\top}$	$\Leftrightarrow \exists c_1, c_2. c \equiv c_1 \otimes c_2$
$e, \sigma, s \models_S \alpha \text{-} @ P_S$	$\Leftrightarrow \exists x, s'. e(\alpha) = x \wedge s' \equiv (x)(s) \wedge e, \sigma, s' \models_S P_S$		$\wedge e, \sigma, c_1 \models_{\top} P_{\top}$
$e, \sigma, s \models_S \forall \alpha. P_S$	$\Leftrightarrow \exists x. x \# e, s \wedge e[\alpha \mapsto x], \sigma, s \models_S P_S$		$\wedge e, \sigma, c_1 \models_{\top} P_{\top}$
$e, \sigma, s \models_S \alpha \leftarrow P_{\top}$	$\Leftrightarrow \exists c, x. e(\alpha) = x \wedge s \equiv x \leftarrow c \wedge e, \sigma, c \models_{\top} P_{\top}$	$e, \sigma, c \models_{\top} c$	$\Leftrightarrow c \equiv e(c)$
$e, \sigma, s \models_S s$	$\Leftrightarrow s \equiv e(s)$	$e, \sigma, c \models_{\top} B$	$\Leftrightarrow \llbracket B \rrbracket \sigma = \mathbf{true}$
$e, \sigma, s \models_S B$	$\Leftrightarrow \llbracket B \rrbracket \sigma = \mathbf{true}$	$e, \sigma, c \models_{\top} @_{\top} \alpha$	$\Leftrightarrow e(\alpha) \in \mathit{free}(c)$

(s)↓ denotes that s is well formed. x#s denotes that x is fresh with respect to s.

We omit the standard semantics for $P \Rightarrow Q$, **false** and $\exists v. P$.

Fig. 6. Satisfaction Relations of Segment Logic for Trees.

Example 6 (Segment Logic Examples).

- (a) The segment formula $\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\delta]$ describes a tree segment consisting of a node n with address α and context hole γ , and node m with address β and context hole δ . The variables n and m cannot denote the same node identifier; similarly, α, β cannot denote the same hole identifier; neither can γ, δ .
- (b) The segment formula $\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\mathit{tree}(c)]$ describes a tree segment consisting of a single node n at address α and a complete tree (a tree with no holes) with top node m at address β . This formula is the safety precondition for the small axiom of the `appendSub(n, m)` command. In particular, the formula states that m cannot be an ancestor of n , as n is disjoint from the tree c .
- (c) The segment formula $H\alpha, \beta. (\delta \leftarrow r[\alpha \otimes \beta] * \alpha \leftarrow n[\gamma] * \beta \leftarrow m[\mathit{tree}(c)])$ describes a tree segment consisting of a node r at address δ whose children are given by the holes α and β , and a tree segment satisfying the formula in Example (b). The labels α, β are under the hiding quantification, and hence denote fresh, unequal, identifiers. This formula is equivalent to $\delta \leftarrow r[n[\gamma] \otimes m[\mathit{tree}(c)]]$, which states that there is a node r at address δ whose children are n and m .
- (d) To specify our language, it is enough to work with the hiding quantification. However, to describe the weakest preconditions, we must use revelation. For example, the weakest precondition of the `deleteSubtree(n)` command is $\exists c. \forall \alpha. \alpha @ ((\alpha \leftarrow n[\emptyset_{\top}] \multimap (\alpha \text{-} @ P_S)) * \alpha \leftarrow n[\mathit{tree}(c)])$. This formula describes a tree segment which can be separated into a complete tree, with top node n at a fresh address x denoted by α , and a segment s satisfying $(\alpha \leftarrow n[\emptyset_{\top}] \multimap (\alpha \text{-} @ P_S))$. The segment s , when extended to $(x)(x \leftarrow n[\emptyset_{\top}] + s)$, satisfies P_S .

5 Local Hoare Reasoning

We use Segment Logic to provide local Hoare reasoning about programs written in the language given in Definition 5. First, we give a fault avoiding, partial correctness interpretation of local Hoare triples following [19]. Informally, $\{P_S\} C \{Q_S\}$ means that, when P_S holds for a tree segment s , then command C does not fault when run on s and the result, if C terminates, satisfies Q_S .

Definition 13 (Local Hoare Triples). Recall the evaluation relation \rightsquigarrow relating configuration triples C, σ, s , terminal states σ, s and faults in **Fig. 4**. The fault-avoiding partial correctness interpretation of local Hoare Triples is given below:

$$\{P_S\} C \{Q_S\} \Leftrightarrow \forall e, \sigma, s. \mathit{free}(C) \subseteq \mathit{dom}(\sigma) \wedge \mathit{free}(P_S) \cup \mathit{free}(Q_S) \subseteq \mathit{dom}(\sigma) \cup \mathit{dom}(e) \\ \wedge e, \sigma, s \models_S P_S \Rightarrow C, \sigma, s \not\rightsquigarrow \mathit{fault} \wedge \forall \sigma', s'. C, \sigma, s \rightsquigarrow \sigma', s' \Rightarrow e, \sigma', s' \models_S Q_S$$

$\{\alpha \leftarrow m[\beta \otimes n[\delta] \otimes \gamma] \wedge (n' = n_0)\}$	$n' := \text{getUp}(n)$	$\{(\alpha \leftarrow m[\beta \otimes n[\delta] \otimes \gamma])_{[n_0/n']} \wedge (n' = m)\}$
$\{\alpha \leftarrow n[\delta] \otimes m[\beta] \wedge (n' = n_0)\}$	$n' := \text{getRight}(n)$	$\{(\alpha \leftarrow n[\delta] \otimes m[\beta])_{[n_0/n']} \wedge (n' = m)\}$
$\{\alpha \leftarrow m[\beta \otimes n[\delta]] \wedge (n' = n_0)\}$	$n' := \text{getRight}(n)$	$\{(\alpha \leftarrow m[\beta \otimes n[\delta]])_{[n_0/n']} \wedge (n' = \text{null})\}$
$\{\alpha \leftarrow n[\delta \otimes m[\beta]] \wedge (n' = n_0)\}$	$n' := \text{getLast}(n)$	$\{(\alpha \leftarrow n[\delta \otimes m[\beta]])_{[n_0/n']} \wedge (n' = m)\}$
$\{\alpha \leftarrow n[\emptyset_T] \wedge (n' = n_0)\}$	$n' := \text{getLast}(n)$	$\{(\alpha \leftarrow n[\emptyset_T])_{[n_0/n']} \wedge (n' = \text{null})\}$
$\{\alpha \leftarrow n[\beta]\}$	$\text{insertNodeAbove}(n)$	$\{\alpha \leftarrow \circ[n[\beta]]\}$
$\{\alpha \leftarrow n[\beta]\}$	$\text{deleteNode}(n)$	$\{\alpha \leftarrow \beta\}$
$\{\alpha \leftarrow n[\text{tree}(c)]\}$	$\text{deleteSubtree}(n)$	$\{\alpha \leftarrow n[\emptyset_T]\}$
$\{\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\delta]\}$	$\text{moveNodeAbove}(n, m)$	$\{\alpha \leftarrow m[n[\gamma]] * \beta \leftarrow \delta\}$
$\{\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\delta]\}$	$\text{appendNode}(n, m)$	$\{\alpha \leftarrow n[\gamma \otimes m[\emptyset_T]] * \beta \leftarrow \delta\}$
$\{\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\text{tree}(c)]\}$	$\text{appendSub}(n, m)$	$\{\alpha \leftarrow n[\gamma \otimes \text{tree}(c)] * \beta \leftarrow m[\emptyset_T]\}$

n and n_0 are distinct. The omitted commands have analogous or standard axioms.

Fig. 7. A Selection of the Small Axioms for the Basic Tree Update Commands.

Definition 14 (Small Axioms). *The Small Axioms for the basic tree update commands from Fig. 4 are given in Fig. 7.*

With Raza, Gardner has developed the formal definitions of footprints and small specifications for abstract local functions using Abstract Separation Logic [15]. It would be interesting to extend this abstract theory to the tree segments and reasoning studied here, and prove that the axioms really are small.

Definition 15 (Inference Rules). *The local reasoning inference rules include the standard Hoare Logic Rules for Sequencing, Consequence, Disjunction, Local Variable, If-Then-Else, While-Do, and the rules for Separation Frame, Revelation Frame, Auxiliary Variable Elimination and Fresh Label Elimination given by:*

$\text{S} \quad \frac{\text{F} \quad \{P_S\} \text{C} \{Q_S\}}{\{P_S * R_S\} \text{C} \{Q_S * R_S\}} \quad \text{mod}(\text{C}) \cap \text{free}(R_S) = \{\}$	$\text{R} \quad \frac{\text{F} \quad \{P_S\} \text{C} \{Q_S\}}{\{\alpha \textcircled{R} P_S\} \text{C} \{\alpha \textcircled{R} Q_S\}}$
$\text{A} \quad \frac{\text{V} \quad \text{E} \quad \{P_S\} \text{C} \{Q_S\}}{\{\exists n. P_S\} \text{C} \{\exists n. Q_S\}} \quad n \notin \text{free}(\text{C})$	$\text{F} \quad \frac{\text{V} \quad \text{E} \quad \{P_S\} \text{C} \{Q_S\}}{\{\forall \alpha. P_S\} \text{C} \{\forall \alpha. Q_S\}}$

Recall that the set of variables modified by a command C is denoted $\text{mod}(\text{C})$.

The Separation Frame rule is standard from Separation Logic and allows us to extend the working tree with tree segments that are not used by the command. The Revelation Frame rule is similar. It allows us to add hiding binders to the working tree, since hole identifiers are not used by any of our commands. The Fresh Variable Elimination rule is analogous to the standard Auxiliary Variable Elimination rule. To see how we use these rules, consider again the small axiom of `appendSub`. We can extend this axiom using our inference rules to obtain:

$$\{\text{H}\alpha, \beta. (\epsilon \leftarrow \text{r}[\alpha \otimes \beta] * \alpha \leftarrow n[\gamma] * \beta \leftarrow m[\text{tree}(c)])\}$$

$$\quad \text{appendSub}(n, m)$$

$$\{\text{H}\alpha, \beta. (\epsilon \leftarrow \text{r}[\alpha \otimes \beta] * \alpha \leftarrow n[\gamma \otimes \text{tree}(c)] * \beta \leftarrow m[\emptyset_T])\}$$

This triple can be simplified using the consequence rule following the discussion in Example 6(c). In this example, it is natural to use the derived hiding quantification. Following the discussion in Example 6(d) we can see that the primitive revelation connectives are, however, necessary to obtain the weakest preconditions, a selection of which are shown in Fig. 8.

$$\begin{array}{l}
\{\exists m, n_0. \text{H}\alpha, \beta, \gamma, \delta. \diamond \alpha \leftarrow m[\beta \otimes n[\delta] \otimes \gamma] \wedge (n' = n_0) \wedge (\alpha, \beta, \gamma, \delta \text{--}\textcircled{R} P_{S[m/n']})\} \quad n' := \text{getUp}(n) \quad \{P_S\} \\
\left\{ \begin{array}{l} \exists m, n_0. \text{H}\alpha, \beta, \delta. \quad \diamond \alpha \leftarrow n[\delta] \otimes m[\beta] \wedge (n' = n_0) \wedge (\alpha, \beta, \delta \text{--}\textcircled{R} P_{S[m/n']}) \\ \vee \diamond \alpha \leftarrow m[\beta \otimes n[\delta]] \wedge (n' = n_0) \wedge (\alpha, \beta, \delta \text{--}\textcircled{R} P_{S[null/n']}) \end{array} \right\} \quad n' := \text{getRight}(n) \quad \{P_S\} \\
\left\{ \begin{array}{l} \exists m, n_0. \text{H}\alpha, \beta, \delta. \quad \diamond \alpha \leftarrow n[\delta \otimes m[\beta]] \wedge (n' = n_0) \wedge (\alpha, \beta, \delta \text{--}\textcircled{R} P_{S[m/n']}) \\ \vee \diamond \alpha \leftarrow n[\emptyset_T] \wedge (n' = n_0) \wedge (\alpha \text{--}\textcircled{R} P_{S[null/n']}) \end{array} \right\} \quad n' := \text{getLast}(n) \quad \{P_S\} \\
\{\text{H}\alpha, \beta. ((\alpha \leftarrow \circ[n[\beta]]) \text{--} * (\alpha, \beta \text{--}\textcircled{R} P_S)) * \alpha \leftarrow n[\beta]\} \quad \text{insertNodeAbove}(n) \quad \{P_S\} \\
\{\text{H}\alpha, \beta. ((\alpha \leftarrow \beta \text{--} * (\alpha, \beta \text{--}\textcircled{R} P_S)) * \alpha \leftarrow n[\beta])\} \quad \text{deleteNode}(n) \quad \{P_S\} \\
\{\exists c. \text{H}\alpha. ((\alpha \leftarrow n[\emptyset_T] \text{--} * (\alpha \text{--}\textcircled{R} P_S)) * \alpha \leftarrow n[\text{tree}(c)])\} \quad \text{deleteSubtree}(n) \quad \{P_S\} \\
\{\text{H}\alpha, \beta, \gamma, \delta. (((\alpha \leftarrow m[n[\gamma]]) * \beta \leftarrow \delta) \text{--} * (\alpha, \beta, \gamma, \delta \text{--}\textcircled{R} P_S)) * (\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\delta])\} \quad \text{moveNodeAbove}(n, m) \quad \{P_S\} \\
\{\text{H}\alpha, \beta, \gamma, \delta. (((\alpha \leftarrow n[\gamma \otimes m[\emptyset_T]]) * \beta \leftarrow \delta) \text{--} * (\alpha, \beta, \gamma, \delta \text{--}\textcircled{R} P_S)) * (\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\delta])\} \quad \text{appendNode}(n, m) \quad \{P_S\} \\
\{\exists c. \text{H}\alpha, \beta, \gamma. (((\alpha \leftarrow n[\gamma \otimes \text{tree}(c)]) * \beta \leftarrow m[\emptyset_T]) \text{--} * (\alpha, \beta, \gamma \text{--}\textcircled{R} P_S)) * (\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\text{tree}(c)])\} \quad \text{appendSub}(n, m) \quad \{P_S\}
\end{array}$$

Fig. 8. A Selection of the Weakest Preconditions for our Basic Tree Update Commands.

Theorem 1 (Soundness and Completeness). *The small axioms and inference rules are sound. For straight line code, they are also complete.*

Proof Sketch. Soundness is straightforward to prove. Completeness for straight line code follows from the derivability of the the weakest preconditions (Fig. 8) from the small axioms (Fig. 7). See the full paper [7] for further details.

Example 7 (Specifying appendChild). In Example 1 we gave the `appendChild` program. The command's specification and its derivation are:

$$\begin{array}{l}
\{\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\text{tree}(c)]\} \\
\text{local } \text{temp} := \text{null} \text{ in } \{ \\
\{\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\text{tree}(c)] \wedge (\text{temp} = \text{null})\} \\
\{\text{H}\delta. \alpha \leftarrow n[\gamma] * \beta \leftarrow m[\delta] * \delta \leftarrow \text{tree}(c) \wedge (\text{temp} = \text{null})\} \\
\quad \text{insertNodeAbove}(m); \\
\{\text{H}\delta. \alpha \leftarrow n[\gamma] * \beta \leftarrow m[\delta] * \delta \leftarrow \text{tree}(c) \wedge (\text{temp} = \text{null})\} \\
\quad \text{temp} := \text{getUp}(m); \\
\{\text{H}\delta. \alpha \leftarrow n[\gamma] * \beta \leftarrow \text{temp}[m[\delta]] * \delta \leftarrow \text{tree}(c)\} \\
\{\text{H}\epsilon, \delta. \alpha \leftarrow n[\gamma] * \beta \leftarrow \text{temp}[\epsilon] * \epsilon \leftarrow m[\delta] * \delta \leftarrow \text{tree}(c)\} \\
\quad \text{appendNode}(n, m); \\
\{\text{H}\epsilon, \delta. \alpha \leftarrow n[\gamma \otimes m[\emptyset_T]] * \beta \leftarrow \text{temp}[\epsilon] * \epsilon \leftarrow \delta * \delta \leftarrow \text{tree}(c)\} \\
\{\alpha \leftarrow n[\gamma \otimes m[\emptyset_T]] * \beta \leftarrow \text{temp}[\text{tree}(c)]\} \\
\{\text{H}\epsilon, \delta. \alpha \leftarrow n[\gamma \otimes \epsilon] * \epsilon \leftarrow m[\delta] * \delta \leftarrow \emptyset_T * \beta \leftarrow \text{temp}[\text{tree}(c)]\} \\
\quad \text{appendSub}(m, \text{temp}); \\
\{\text{H}\epsilon, \delta. \alpha \leftarrow n[\gamma \otimes \epsilon] * \epsilon \leftarrow m[\delta \otimes \text{tree}(c)] * \delta \leftarrow \emptyset_T * \beta \leftarrow \text{temp}[\emptyset_T]\} \\
\{\alpha \leftarrow n[\gamma \otimes m[\text{tree}(c)]] * \beta \leftarrow \text{temp}[\emptyset_T]\} \\
\quad \text{deleteNode}(\text{temp}) \quad \} \\
\{\alpha \leftarrow n[\gamma \otimes m[\text{tree}(c)]] * \beta \leftarrow \emptyset_T\}
\end{array}$$

Throughout the proof, we use the rules to separate out the footprint of a command, apply the appropriate small axiom, then use the rules to compress the result back into the original tree.

Example 8 (Specifying Node Manipulation). In Examples 2, 3 and 4 we gave three node manipulation programs: `simple(n)`, `nodeSwap(n, m)` and `nodeCycle(n)`. The specifications for each of these programs are:

$$\begin{array}{lll}
\{\alpha \leftarrow n[m[\beta] \otimes \gamma]\} & \{\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\delta]\} & \{\alpha \leftarrow n[m[\beta] \otimes \gamma \otimes 1[\delta]]\} \\
\text{simple}(n) & \text{nodeSwap}(n, m) & \text{nodeCycle}(n) \\
\{\alpha \leftarrow m[n[\beta] \otimes \gamma]\} & \{\alpha \leftarrow m[\gamma] * \beta \leftarrow n[\delta]\} & \{\alpha \leftarrow 1[n[\beta] \otimes \gamma \otimes m[\delta]]\}
\end{array}$$

The derivations of these specifications are shown in Fig. 9.

Example 9 (Specifying queuePop). In a similar way, we can derive the following specification for the `queuePop` program from Example 5:

$$\begin{array}{l}
\{\alpha \leftarrow n[m[\text{tree}(c)] \otimes \gamma \otimes i[\beta]]\} \\
\quad \text{queuePop}(n) \\
\{\alpha \leftarrow m[\gamma \otimes n[\beta] \otimes i[\text{tree}(c)]]\}
\end{array}$$

simple derivation:	nodeSwap derivation:	nodeCycle derivation:
$\{\alpha \leftarrow n[m[\beta] \otimes \gamma]\}$ $\text{local } temp, first := null \text{ in } \{$ $\{\alpha \leftarrow n[m[\beta] \otimes \gamma] \wedge (temp = null)\}$ $\wedge (first = null)\}$ $\text{insertNodeAbove}(n);$ $\{\alpha \leftarrow o[n[m[\beta] \otimes \gamma] \wedge (temp = null)\}$ $\wedge (first = null)\}$ $temp := getUp(n);$ $\{\alpha \leftarrow temp[n[m[\beta] \otimes \gamma] \wedge (first = null)\}$ $first := getFirst(n);$ $\{\alpha \leftarrow temp[n[first[\beta] \otimes \gamma] \wedge (first = m)\}$ $\text{moveNodeAbove}(first, n);$ $\{\alpha \leftarrow temp[n[first[\beta] \otimes \gamma] \wedge (first = m)\}$ $\text{moveNodeAbove}(temp, first);$ $\{\alpha \leftarrow first[temp[n[\beta] \otimes \gamma] \wedge (first = m)\}$ $\text{deleteNode}(temp)\}$ $\{\alpha \leftarrow first[n[\beta] \otimes \gamma] \wedge (first = m)\}$ $\{\alpha \leftarrow m[n[\beta] \otimes \gamma]\}$	$\{\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\delta]\}$ $\text{local } temp := null \text{ in } \{$ $\{\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\delta]\}$ $\wedge (temp = null)\}$ $\text{insertNodeAbove}(n);$ $\{\alpha \leftarrow o[n[\gamma] * \beta \leftarrow m[\delta]\}$ $\wedge (temp = null)\}$ $temp := getUp(n);$ $\{\alpha \leftarrow temp[n[\gamma] * \beta \leftarrow m[\delta]\}$ $\text{moveNodeAbove}(m, n);$ $\{\alpha \leftarrow temp[\gamma] * \beta \leftarrow n[m[\delta]]\}$ $\text{moveNodeAbove}(temp, m);$ $\{\alpha \leftarrow m[temp[\gamma] * \beta \leftarrow n[\delta]]\}$ $\text{deleteNode}(temp)\}$ $\{\alpha \leftarrow m[\gamma] * \beta \leftarrow n[\delta]\}$	$\{\alpha \leftarrow n[m[\beta] \otimes \gamma \otimes l[\delta]]\}$ $\text{local } first, last := null \text{ in } \{$ $\{\alpha \leftarrow n[m[\beta] \otimes \gamma \otimes l[\delta]]\}$ $\wedge (first = null) \wedge (last = null)\}$ $first := getFirst(n);$ $\{\alpha \leftarrow n[first[\beta] \otimes \gamma \otimes l[\delta]]\}$ $\wedge (first = m) \wedge (last = null)\}$ $last := getLast(n);$ $\{\alpha \leftarrow n[first[\beta] \otimes \gamma \otimes last[\delta]]\}$ $\wedge (first = m) \wedge (last = l)\}$ $\text{nodeSwap}(n, last);$ $\{\alpha \leftarrow last[first[\beta] \otimes \gamma \otimes n[\delta]]\}$ $\wedge (first = m) \wedge (last = l)\}$ $\text{nodeSwap}(n, first)\}$ $\{\alpha \leftarrow last[n[\beta] \otimes \gamma \otimes first[\delta]]\}$ $\wedge (first = m) \wedge (last = l)\}$ $\{\alpha \leftarrow l[n[\beta] \otimes \gamma \otimes m[\delta]]\}$

Fig. 9. Derivations of the Specifications for simple, nodeSwap and nodeCycle.

6 Conclusion

We have introduced Segment Logic for reasoning about structured data update in general, and tree update in particular. Using Segment Logic, we have demonstrated that it is possible to give small axioms for tree update commands such as DOM's `appendChild` command. In this paper, we have concentrated on a simple, lightweight tree update language. It is straightforward to transfer the techniques developed here to Featherweight DOM [6]. We do not envisage difficulties with extending the approach to the full DOM specification [16].

A typical Segment Logic proof separates the working tree into segments, identifying the tree segment which corresponds to the footprint of a command. It applies a small axiom to this segment, and compresses the updated fragment back into the original tree. This separation and compression is key to our Segment Logic reasoning. It is not unlike the unfolding and folding of abstract predicates, due to Parkinson and Vafeiadis [17]. The difference is that reasoning using Separation Logic with abstract predicates is implementation dependent: the formula $slist(l, i)$ describes a list l implemented as a singly-linked list with heap address i ; the formula $dlist(l, i, j)$ describes a list l implemented as a doubly-linked list with heap addresses i and j . By contrast, reasoning using Segment Logic is implementation independent: the formula $\alpha \leftarrow P$ describes e.g. a list or tree identified and satisfying formula or data type P at abstract address α .

Our next step is to design and formally specify a concurrent XML update language, combining ideas from Featherweight DOM, Concurrent Separation Logic [14] and Segment Logic. We believe that Segment Logic provides us with crucial technology for achieving this goal. For example, consider the program `deleteTree(n) || deleteTree(m)`, which should succeed if the two trees being called are disjoint. A Segment Logic specification of this program is:

$$\{\alpha \leftarrow n[\text{tree}(c_1)] * \beta \leftarrow m[\text{tree}(c_2)]\}$$

$$\text{deleteTree}(n) \parallel \text{deleteTree}(m)$$

$$\{\alpha \leftarrow \emptyset_T * \beta \leftarrow \emptyset_T\}$$

Segment Logic allows us to establish such natural disjointness properties, since it combines reasoning directly about the abstract tree structure with using the

separating conjunction $*$. Our goal is to extend the update language presented here with parallel composition and critical regions, and adapt the Concurrent Separation Logic reasoning to provide a formal, compositional specification of a concurrent XML update language.

Acknowledgments: We thank Thomas Dinsdale-Young for many interesting discussions regarding this work and Viktor Vafeiadis for the name Segment Logic. Gardner acknowledges support of a Microsoft/RAEng Senior Research Fellowship. Wheelhouse acknowledges support of an EPSRC DTA award.

References

1. J. Berdine, C. Calcagno, and P. W. O’Hearn. Smallfoot: Modular automatic assertion checking with separation logic. *FMCO, LNCS*, 4111, 2005.
2. C. Calcagno, T. Dinsdale-Young, and P. Gardner. Adjunct elimination in context logic for trees. *APLAS, LNCS*, 4807, 2007.
3. C. Calcagno, P. Gardner, and U. Zarfaty. Context logic and tree update. *POPL, ACM*, 40, 2005.
4. L. Cardelli and A. D. Gordon. Ambient logic. *Mathematical Structures in Computer Science, in press*, 2006.
5. M. J. Gabbay and A. M. Pitts. A new approach to abstract syntax with variable binding. *Formal Aspects of Computing*, 13, 2002.
6. P. Gardner, G. Smith, M. Wheelhouse, and U. Zarfaty. Local Hoare reasoning about DOM. *PODS, ACM*, 27, 2008.
7. P. Gardner and M. Wheelhouse. Small specifications for tree update, extended version. <http://www.doc.ic.ac.uk/~mjlw03/PersonalWebpage/pdfs/moveFull.pdf>, 2009.
8. P. Gardner and U. Zarfaty. Reasoning about high-level tree update and its low-level implementation. Technical Report DTR09-9, Imperial College, 2009.
9. H. Hosoya and B. C. Pierce. Xduce: A statically typed XML processing language. *TOIT, ACM*, 3, 2003.
10. S. Ishtiaq and P. W. O’Hearn. BI as an assertion language for mutable data structures. *POPL, ACM*, 36, 2001.
11. R. Milner. Pi-nets: A graphical form of π -calculus. *ESOP, LNCS*, 788, 1994.
12. R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes I & II. *Information and Computation*, 100, 1992.
13. P. W. O’Hearn, J. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. *CSL, LNCS*, 15, 2001.
14. P. W. O’Hearn. Resources, concurrency and local reasoning. *Theoretical Computer Science*, 375, 2007.
15. M. Raza and P. Gardner. Footprints in local reasoning. *FoSSaCS, LNCS*, 4962, 2008.
16. G. Smith. Providing a formal specification for DOM core level 1. PhD Thesis, to be submitted December 2009.
17. V. Vafeiadis. Modular fine-grained concurrency verification. Technical Report UCAM-CL-TR-726, Cambridge, 2008.
18. W3C. Dom: Document object model. W3C recommendation. <http://www.w3.org/DOM/>, 2005.
19. H. Yang and P. W. O’Hearn. A semantic basis for local reasoning. *FoSSaCS, LNCS*, 2303, 2002.