

Small Specifications for Tree Update

Philippa Gardner and Mark Wheelhouse

Imperial College London, {pg, mjlw03}@doc.ic.ac.uk

Abstract. O’Hearn, Reynolds and Yang introduced local Hoare reasoning about mutable data structures using Separation Logic. They reason about the local parts of the memory accessed by programs, and thus construct their smallest complete specifications. Gardner *et al.* generalised their work, using Context Logic to reason about structured data at the same level of abstraction as the data itself. In particular, they developed a formal specification of the Document Object Model, a W3C XML update library. Whilst they kept to the spirit of local reasoning, they were not able to retain small specifications: for example, the specification of `appendChild` was not small. We show how to obtain small specifications by working with a more fine-grained context structure, allowing us to work with arbitrary tree fragments.

Key words: Specification, logical reasoning, program verification, locality

1 Introduction

Separation Logic [12], 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 the memory (the footprint) that are accessed by a program. In particular, they introduced *small axioms* for specifying the atomic commands, using the smallest heaps possible to obtain complete specifications of programs, and the *frame rule* for extending the reasoning to larger heaps. The resulting modular reasoning has been used to notable success for verifying memory safety properties of large C-programs.

Calcagno, Gardner and Zarfaty generalised Separation Logic to reason about more complex data structures, such as those found on the web, by providing a fundamental shift in the reasoning. Structured data update typically identifies the portion of data to be replaced, removes it, and inserts the new data in the same place. Gardner *et al.* introduced Context Logic to reason about both data and this place of insertion (contexts). Their original work applied Context Logic to reason about a simple tree update language, with analogous small axioms for the basic tree update commands and a generalised frame rule.

With Smith and Zarfaty, Gardner and Wheelhouse have applied Context Logic to provide a concise, compositional specification of the W3C Document Object Model (DOM) [17],[7], a library for XML update. In our initial paper[6], we introduced and reasoned about Featherweight DOM (called Minimal DOM in the paper), a fragment of DOM which concentrates on the DOM tree structure rather than the full DOM structure. The compositionality of our reasoning means that, as well as specifying the basic DOM commands, we can also reason about simple JavaScript programs which call DOM. Gardner and Smith have extended




Fig. 1. Splitting up the working tree using multi-holed contexts.

the reasoning to the full DOM Core Level 1 specification[15]. This extension was a substantial piece of work, not because of the reasoning, but because the full DOM specification is large, underspecified and difficult to interpret.

Context Logic reasoning can be adapted to many familiar context styles associated with structured data. In our initial papers and the DOM work, we chose to use single-holed contexts because it was enough to introduce our ideas of reasoning about tree update. However, although our DOM specification keeps to the spirit of local reasoning, it does not have small axioms for all the atomic commands. This can be illustrated by the command `appendChild(n, m)` which moves the tree with top node identified by DOM identifier m to be the last child of the tree identified by n . Since n and m may be in distinct parts of the tree, it certainly seems natural to move to multi-holed Context Logic[2]. However, we shall see that multi-holed contexts are not enough.

Consider Figure 1 which 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)` in multi-holed Context Logic is:

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

Figure 1 shows, in each successful case, how the tree satisfies the precondition. The precondition specifies that the working tree can be split into a subtree with top node identified by m , and a context with hole variable β (equals y in the figure) satisfying context formula $C \circ_{\alpha} n[c_1]$. This formula $C \circ_{\alpha} n[c_1]$ states that the context can be split into a subcontext with top node n and an unspecified context with hole α (equals x in figure) given by context variable C . The postcondition states that the tree at m moves to be the last child of n , the empty tree replaces the tree at m , and the surrounding context denoted by variable C remains the same.

The problem with this `appendChild(n, m)` axiom is that it is not small, since it uses variable C to stand for a surrounding context which contains both n and m . We could put additional constraints on C to insist that the context is minimal, but this is not the point. Intuitively, 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 . We need a finer way of splitting the tree to be able to capture this footprint.




Fig. 2. Splitting up the working tree using tree fragments.

We use *tree fragments*. The idea is that the working tree can be split into tree fragments that can be reasoned about separately, but which still ‘know’ how they join back together. They are similar to multi-holed tree contexts in that they have unique hole labels; they are different in the way they join together. Tree fragments have unique hole addresses, which determine which holes the fragments can fill. With multi-holed contexts, it is the application function that determines which hole gets filled. Consider figure 2. In both example cases, the working tree is split into a bunch of tree fragments; the hole labels and addresses determine how the tree fragments join back up to form the original tree.

There are several further features about our tree fragments to observe. Consider the right-hand equalities of figure 2. In both cases, the tree fragment with top node n has been split into just the top node n with the same address and fresh hole label z , and another tree fragment with address z . 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`. If we take this to the extreme, we can cut up the tree structure into a collection of nodes with hole spaghetti (similar to heap cells), where the hole labels and addresses show how the nodes are joined together. Although this is possible, this is not how we use the hole labels and addresses. We only cut up the tree in a minimal way in order to provide the right fragment about which to reason. Also notice that the hole labels and addresses have been hidden by a freshness operator; the ν in the figure. With the restriction, the fragments can be compressed into the larger fragments indicated by the figure. Without the restriction, the fragments cannot be compressed and the hole labels and addresses would behave rather like pointers and addresses in a heap.

We introduce Context Logic for analysing tree fragments. It is analogous to our previous work on single-holed and multi-holed Context Logic in the sense that we analyse fragments of high-level trees. It is different from the previous work in that we use the commutative separating conjunction $*$ of Separation Logic rather than a non-commutative separating application. We also use the revelation connectives and freshness quantification of Gabbay and Pitts [5] and Cardelli and Gordon [4]. Interestingly, we shall see that these constructs are important for the weakest preconditions. Using this Context Logic for analysing tree fragments, we are able to give a small axiom for `appendChild`(n, m):

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

The precondition specifies two tree fragments: a tree fragment at variable address α with node n and a tree fragment at address β with a complete tree whose top node is m . The postcondition states that the tree at m moves to be the last child of n being replaced by the empty tree. The axiom is small, in the sense that it captures the intuitive footprint of `appendChild(n, m)`. We can extend the axiom to larger tree fragments using the normal frame rule for separation conjunction, a non-standard frame rule for revelation, and a rule for freshness quantification.

We must especially point out the difference in spirit between this work on reasoning about high-level tree update, and the work of O’Hearn, Parkinson and colleagues on reasoning about mutable data structures represented in heaps. The reasoning is at a different level of abstraction. O’Hearn and Parkinson are working with C-programs and object-oriented programs, where it is natural to work with the basic heap model and build up layers of abstraction. Our focus is on high-level tree update languages such as that specified by the DOM library, where we must work with the tree structures directly.

2 Tree Update Language

We present a simple, but expressive, high-level tree update language. Our tree structures are left intentionally simple. We work with finite, ordered, unranked trees and tree contexts [2], with unique node identifiers for specifying the locations of updates as in DOM. It is straightforward to incorporate (and reason about) additional data such as tag information and text data. 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 labels.

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

$$\begin{array}{ll} \text{tree context } c ::= & \emptyset_C \quad \text{empty tree} \\ & x \quad \text{tree context hole label} \\ & n[c] \quad \text{tree context with top node } n \\ & c \otimes c \quad \text{tree composition} \end{array}$$

with the restriction that each hole label, $x \in X$, 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 the \otimes operator is associative with identity \emptyset_C . The set of hole labels that occur in tree context c is denoted by $fn(c)$. We use t, t_1, t_2 to denote tree contexts with no context holes.

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 :

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

We abbreviate $ap_x(c_1, c_2)$ by $c_1 \circledast c_2$. 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, imperative language, based on variable assignment, and update commands. The program state is made up of

two components. The first component is the working tree which contains all of the nodes we will be manipulating with our programs. The second component is a high-level variable store containing variables for both node identifiers and tree-shapes (trees modulo renaming of identifiers, allowing high-level manipulation of tree structures). The choice of having tree shapes, as opposed to trees, in the store illustrates a seemingly paradoxical property of high-level, imperative update: while some way of identifying nodes is required to specify the location of in-place updates, these identifiers are typically not considered an important part of the high-level structure itself.

Definition 3 (Tree-shapes). Tree-shapes $t_o \in T_o$ are defined by the grammar:

$$\begin{aligned} \text{tree-shape } t_o ::= & \quad \emptyset_T \quad \text{empty tree} \\ & \quad \circ[t_o] \quad \text{tree node} \\ & \quad t_o \otimes t_o \quad \text{tree composition} \end{aligned}$$

We write $\langle t \rangle$ for the shape of a tree t , where $\langle \emptyset_C \rangle = \emptyset_T$, $\langle n[t] \rangle = \circ[\langle t \rangle]$ and $\langle t \otimes t' \rangle = \langle t \rangle \otimes \langle t' \rangle$. We only store complete trees. We write $t \simeq t'$ when $\langle t \rangle = \langle t' \rangle$.

Definition 4 (Variable Store). The variable store $s \in S$ consists of a pair of finite partial functions

$$s : (Var_I \rightarrow_{fin} I \cup \{\mathbf{null}\}) \times (Var_{T_o} \rightarrow_{fin} T_o)$$

mapping location name variables $Var_I = \{m, n, \dots\}$ to location names or \mathbf{null} , and tree shape variables $Var_{T_o} = \{t, \dots\}$ to tree shapes. We write $s[n \mapsto n]$ for the variable store s overwritten with $s(n) = n$, and similarly for $s[t \mapsto t_o]$.

To specify location name and tree-shape values, our language uses simple expressions. Location names are specified either with location name variables or the constant \mathbf{null} ; we forbid direct reference to constant location names other than \mathbf{null} . Tree-shapes are specified using a combination of tree-shape variables and constant tree-shape structures. We also require simple Boolean expressions.

Definition 5 (Expressions). Location name expressions $N \in Exp_I$, tree-shape expressions $T \in Exp_{T_o}$ and Boolean expressions $B \in Exp_B$ are defined by the grammars:

$$\begin{aligned} N ::= & n \mid \mathbf{null} & n \in Var_I \\ T ::= & \emptyset_T \mid t \mid \circ[T] \mid T \otimes T & t \in Var_{T_o} \\ B ::= & N = N \mid T = T \mid \mathbf{false} \mid B \Rightarrow B \end{aligned}$$

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

In previous work [9], we concentrated mainly on tree update commands for changing a tree with some top node n . Here, we give node commands and subtree commands for changing the subtree under node n . The language here is more flexible, allowing us to manipulate pieces of trees; the language in [3] is implementable. The node update commands consist of look-up commands that return a neighboring node in the tree, a delete command that removes a node from the tree, one node insertion command that puts a fresh node into

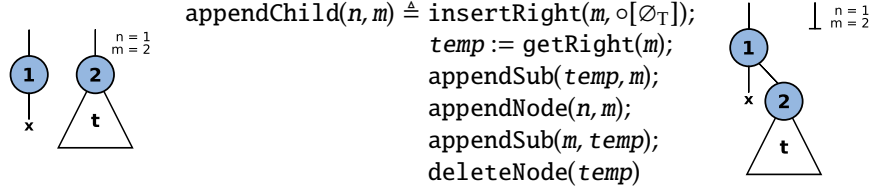
the tree (the others are derivable), and node move commands that take a node out of the tree and replace it in a new position. The node movement commands leave the children of the moved node m as children of m 's old parent. The tree update commands work on subtrees of an identified node and consist of a copy command that stores the shape of a subtree, a delete command that removes an entire subtree from the tree, insertion commands that add new nodes to the tree and subtree move commands that take a subtree out of the tree and replace it in a new position.

Definition 6 (Tree Update Language). *The commands of the tree update language are defined by the node update commands $\mathbf{C}_{\text{nodeUp}}$, the tree update commands $\mathbf{C}_{\text{treeUp}}$, and the standard skip, variable assignment, sequencing, if-then-else and while commands:*

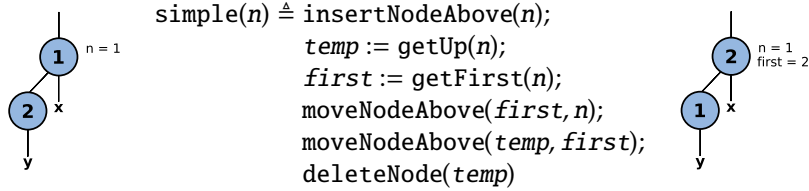
$\mathbf{C}_{\text{nodeUp}} ::= n' := \text{getUp}(n)$	get parent of node n
$n' := \text{getLeft}(n)$	get previous sibling of node n
$n' := \text{getRight}(n)$	get next sibling of node n
$n' := \text{getFirst}(n)$	get first child of node n
$n' := \text{getLast}(n)$	get last child of node n
$\text{deleteNode}(n)$	delete node n
$\text{insertNodeAbove}(n)$	insert a new node above node n
$\text{moveNodeAbove}(n, m)$	move node m above node n
$\text{moveNodeLeft}(n, m)$	move node m to the left of node n
$\text{moveNodeRight}(n, m)$	move node m to the right of node n
$\text{prependNode}(n, m)$	prepend node m to children of node n
$\text{appendNode}(n, m)$	append node m to children of node n
$\mathbf{C}_{\text{treeUp}} ::= x := \text{copy}(n)$	copy shape of subtree starting at node n
$\text{deleteSubtree}(n)$	delete subtree beneath node n
$\text{insertLeft}(n, T)$	insert tree shape T to the left of node n
$\text{insertRight}(n, T)$	insert tree shape T to the right of node n
$\text{insertFirst}(n, T)$	insert tree shape T as first child of n
$\text{insertLast}(n, T)$	insert tree shape T as last child of n
$\text{moveSubLeft}(n, m)$	move children of node m to the left of node n
$\text{moveSubRight}(n, m)$	move children of node m to the right of node n
$\text{prependSub}(n, m)$	prepend children of node m to children of node n
$\text{appendSub}(n, m)$	append children of node m to children of node n

The intuitive behavior of these commands should be self-explanatory. These commands are sufficient to express a wide range of tree manipulation. For example, allocation of a new tree or node can be expressed by the insertion of a literal tree-shape. In particular, inserting the tree-shape expression $\circ[\emptyset_T]$ creates a single new node, with fresh identifier, at a location given by the insert command. The only node insertion command that we need to give explicitly is for inserting a fresh node above an existing node. Our command set is not minimal, for example we could derive the copy command using a combination of lookup, insertion and recursion. We believe the commands chosen lead to a natural and expressive tree update language. We give the operational semantics in section 3, using tree fragments rather than trees, as it simplifies the interpretation of the Hoare triples in section 5.

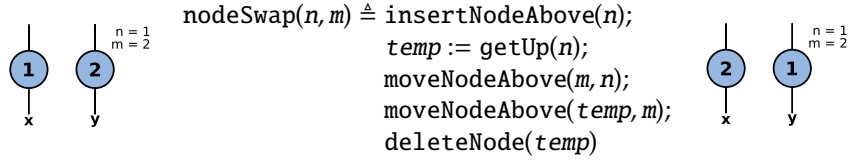
Example 1 (Move). DOM uses the command `appendChild`, whereas here we have `appendNode` and `appendSub`. We use our basic commands to provide the standard `appendChild(n, m)` command. The diagrams illustrate the effect of the program on the part of a tree necessary for the program to run without faulting, with $n = 1$ and $m = 2$. The complete subtree beneath node m is needed as the whole tree is cut out of its original place in the tree and appended to node n .




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




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:



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



Example 5 (Move and Node Swap). Consider a simple hierarchical queuing system given by the program `queuePop(n)` which puts the top element of the hierarchy to the back of the queue and promotes the next element to the top of the queue, carefully maintaining the data related to these elements:



In Section 6 we look at the reasoning of these programs and show how we can specify their behavior from the specifications of their component commands.

3 Tree Fragments

We now give our definition of tree fragments.

Definition 7 (Tree Fragments). Tree fragments $f \in F_{\mathbf{I}, \mathbf{X}}$ are defined by the grammar:

$$\begin{aligned} \text{tree fragment } f ::= & \quad \emptyset_F && \text{empty tree fragment} \\ & x \leftarrow c && \text{tree context } c \text{ with hole address } x \\ & f + f && \text{disjoint union} \\ & (\nu x)(f) && \text{label restriction} \end{aligned}$$


with the restriction that each label, $x \in \mathbf{X}$, occurs free at most once as a hole address and at most once as a hole label in a tree fragment f , and each location name, $n \in \mathbf{I}$, occurs at most once in the a fragment f . Tree fragments are also required to be cycle free. The set of hole labels and hole addresses that occur free in tree fragment f is denoted $fn(f)$.

Definition 8 (Tree Fragment Equivalence). An equivalence relation \equiv over tree fragments is defined by the following axioms:

$$\begin{aligned} f + \emptyset_F &\equiv f \\ f_1 + f_2 &\equiv f_2 + f_1 \\ f_1 + (f_2 + f_3) &\equiv (f_1 + f_2) + f_3 \\ (\nu x)(\emptyset_F) &\equiv \emptyset_F \\ (\nu x)(\nu y)(f) &\equiv (\nu y)(\nu x)(f) \\ (\nu x)(f) &\equiv (\nu y)(f[y/x]) && \text{if } y \notin fn(f) \\ (\nu x)(y \leftarrow c + f) &\equiv y \leftarrow c + (\nu x)(f) && \text{if } x \neq y \text{ and } x \notin fn(c) \\ (\nu x)(y \leftarrow c_1 + x \leftarrow c_2) &\equiv y \leftarrow c_1 \oplus c_2 && \text{if } x \in fn(c_1) \end{aligned}$$

Most of these axioms involving restriction are unsurprising and follow from the π calculus [11]. The last restriction axiom is crucial and enables us to split and join tree fragments at will, as illustrated in Figure 2 of the introduction.

Restriction is well known as a mechanism for hiding names (wires) in Milner's process graphs in particular, and in arbitrary graphs in general. Our tree fragments have similarities and differences with this approach. The tree fragment $(\nu y)(z \leftarrow 1[x \otimes y] + x \leftarrow 2 + y \leftarrow 3)$ can be illustrated as:



$$\begin{array}{c}
\frac{s(n) = n \quad f \equiv (\nu x)(f' + x \leftarrow n[t])}{t := \text{copy}(n), s, f \rightsquigarrow s[t \mapsto \langle n[t] \rangle], f} \quad \frac{s(n) = n \quad f \equiv (\nu w, x, y, z)(f' + x \leftarrow m[y \otimes n[w] \otimes z])}{n' := \text{getUp}(n), s, f \rightsquigarrow s[n' \mapsto m], f} \\
\\
\frac{s(n) = n \quad f \equiv (\nu x, y, z)(f' + x \leftarrow n[y] \otimes m[z])}{n' := \text{getRight}(n), s, f \rightsquigarrow s[n' \mapsto m], f} \quad \frac{s(n) = n \quad f \equiv (\nu x, y, z)(f' + x \leftarrow m[z \otimes n[y]])}{n' := \text{getRight}(n), s, f \rightsquigarrow s[n' \mapsto \text{null}], f} \\
\\
\frac{s(n) = n \quad f \equiv (\nu x, y, z)(f' + x \leftarrow n[y \otimes m[z]])}{n' := \text{getLast}(n), s, f \rightsquigarrow s[n' \mapsto m], f} \quad \frac{s(n) = n \quad f \equiv (\nu x, y, z)(f' + x \leftarrow n[\otimes_C])}{n' := \text{getLast}(n), s, f \rightsquigarrow s[n' \mapsto \text{null}], f} \\
\\
\frac{s(n) = n \quad f \equiv (\nu x, y)(f'' + x \leftarrow n[y]) \quad f' \equiv (\nu x, y)(f'' + x \leftarrow y)}{\text{deleteNode}(n), s, f \rightsquigarrow s, f'} \quad \frac{s(n) = n \quad f \equiv (\nu x)(f'' + x \leftarrow n[t]) \quad f' \equiv (\nu x)(f'' + x \leftarrow n[\otimes_C])}{\text{deleteSubtree}(n), s, f \rightsquigarrow s, f'} \quad \frac{s(n) = n \quad \langle t \rangle = \llbracket T \rrbracket_s \quad t \text{ has fresh ids} \quad f \equiv (\nu x, y)(f'' + x \leftarrow n[y]) \quad f' \equiv (\nu x, y)(f'' + x \leftarrow n[y] \otimes t)}{\text{insertRight}(n, T), s, f \rightsquigarrow s, f'} \\
\\
\frac{s(n) = n \quad f \equiv (\nu x, y)(f'' + x \leftarrow n[y]) \quad n' \text{ fresh id} \quad f' \equiv (\nu x, y)(f'' + x \leftarrow n'[n[y]])}{\text{insertNodeAbove}(n), s, f \rightsquigarrow s, f'} \quad \frac{s(n) = n \quad f \equiv (\nu w, x, y, z)(f'' + x \leftarrow n[z] + y \leftarrow m[w]) \quad s(m) = m \quad f' \equiv (\nu w, x, y, z)(f'' + x \leftarrow m[n[z]] + y \leftarrow w)}{\text{moveNodeAbove}(n, m), s, f \rightsquigarrow s, f'} \\
\\
\frac{s(n) = n \quad f \equiv (\nu w, x, y, z)(f'' + x \leftarrow n[z] + y \leftarrow m[w]) \quad s(m) = m \quad f' \equiv (\nu w, x, y, z)(f'' + x \leftarrow n[z \otimes m] + y \leftarrow w)}{\text{appendNode}(n, m), s, f \rightsquigarrow s, f'} \quad \frac{s(n) = n \quad f \equiv (\nu x, y, z)(f'' + x \leftarrow n[z] + y \leftarrow m[t]) \quad s(m) = m \quad f' \equiv (\nu x, y, z)(f'' + x \leftarrow n[z \otimes t] + y \leftarrow m)}{\text{appendSub}(n, m), s, f \rightsquigarrow s, f'}
\end{array}$$

The cases for skip, assignment, sequencing, if-then-else and while-do are omitted as they are standard. For get, insert, and move only some of the cases are given; the other cases are analogous. Our commands fault when the datastructure does not satisfy any of the preconditions for that command.

Fig. 3. Operational Semantics of the Tree Update Language

and is analogous to the graph approach. However, we are not only using hole labels for wires. Consider the `appendChild` command in example 1. The tree fragment $(z \leftarrow 1[x] + y \leftarrow 2[t])$ updates to $(z \leftarrow 1[x \otimes 2[t]] + y \leftarrow \otimes_C)$: before update the fragment $y \leftarrow 2[t]$ states that a tree can be put in hole z ; after update $y \leftarrow \otimes_C$ states that the empty tree can be put on hole z . In general, unlike heaps, we do not have a sense of arity being preserved by update: before update a node can have a certain number of children; after update it can have a different number of children. The closest work to tree fragments that we have come across is work by Back, which does not have restriction, but otherwise is analogous.

Our programming language manipulates nodes and complete trees. It does not refer to hole labels or hole addresses in any way. However, the operational semantics are greatly simplified by using either tree contexts or tree fragments. We use tree fragments, as this leads to a simpler interpretation of Hoare triples in section 5. We give the operational semantics of the tree update language in Figure 3. We use an evaluation relation \rightsquigarrow relating configuration triples \mathbb{C}, s, f , terminal states s, f , and faults, where f refers to a tree fragment and the free program variables of a command \mathbb{C} are $\text{free}(\mathbb{C})$.

Our style of reasoning 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. A state can be made larger via disjoint tree fragment union or via label restriction. Separate from the reasoning, we also believe that the property of locality leads to good language design. Low-level imperative commands are typically local [12]. In our work on specifying

DOM [7] we demonstrated that the DOM commands are also local. Here, we insist on locality. Consider for example 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 context 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 if it is to be local. The behavior of the other update cases are similar.

4 Context Logic

First, we present the *logical environment* which is a set of functions mapping logical tree context variables to tree contexts, tree fragment variables to tree fragments, and context label variables to context labels. These variables allow us to refer to unchanged data in our pre- and post-conditions (see Definition 14), and make use of quantification in our weakest preconditions (see Figure 6). Tree-shape program variables refer to specific store values and are hence not quantified. We permit location name variables to have the standard dual role as both program variables and logical variables, hence they can be quantified.

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

$$e : (LVar_C \rightarrow C_{I,X}) \times (LVar_F \rightarrow F_{I,X}) \times (LVar_X \rightarrow X)$$

mapping tree context variables $LVar_C = \{c, \dots\}$ to tree contexts, tree fragment variables $LVar_F = \{f, \dots\}$ to tree fragments and label variables $LVar_X = \{\alpha, \beta, \gamma, \delta, \dots\}$ to labels.

We write $e[x \mapsto v]$ for the environment e overwritten with $e(x) = v$.

We are going to work with Context Logic for tree fragments. In fact, our logic has much in common with Separation Logic [12]. In particular, we have the standard classical formulae (additive connectives) and structural formulae (multiplicative connectives) from Separation Logic. The most important of these are the separation connective $*$ and its right adjoint \multimap . Given tree fragment formulae P_F and P'_F : the formula $P_F * P'_F$ describes a tree fragment that can be split into a tree fragment satisfying P_F and a separate tree fragment satisfying P'_F ; and the formula $P_F \multimap P'_F$ describes a tree fragment which, when joined to a tree fragment satisfying P_F , results in a tree fragment satisfying P'_F .

We include label restriction in our tree fragments, which means it is natural to have freshness quantification $\mathcal{I}\alpha$ ¹ and revelation connectives \textcircled{R} and $\neg\textcircled{R}$ from Ambient Logic [5],[4]. These constructs are essential for our weakest preconditions. Given tree fragment formula P_F and hole label α : the formula $\mathcal{I}\alpha. P_F$ describes a tree fragment that with a fresh label stored in variable α satisfies P_F ; the formula $\alpha\textcircled{R}P_F$ describes a tree fragment with a top level restriction of the value of α and, after removing that restriction, the remaining tree fragment satisfies P_F ; and the formula $\alpha\neg\textcircled{R}P_F$ describes a tree fragment which satisfies P_F once it has been extended with a restriction over label stored in variable α .

¹ For our model, it would be possible to use the existential quantification for hole labels instead of the freshness quantification. We choose freshness since it is the natural quantification to accompany revelation.

$e, s, c \models_C P_C \Rightarrow P'_C \Leftrightarrow e, s, c \models_C P_C$	$e, s, f \models_F P_F \Rightarrow P'_F \Leftrightarrow e, s, f \models_F P_F \Rightarrow e, s, f \models_F P'_F$
$e, s, c \models_C \text{false}_C \Leftrightarrow \text{never}$	$e, s, f \models_F \text{false}_F \Leftrightarrow \text{never}$
$e, s, c \models_C \emptyset_C \Leftrightarrow c \equiv \emptyset_C$	$e, s, f \models_F \emptyset_F \Leftrightarrow f \equiv \emptyset_F$
$e, s, c \models_C \alpha \Leftrightarrow c \equiv e(\alpha)$	$e, s, f \models_F \alpha \leftarrow P_C \Leftrightarrow \exists c, x. e(\alpha) = x \wedge f \equiv x \leftarrow c \wedge e, s, c \models_C P_C$
$e, s, c \models_C n[P_C] \Leftrightarrow \exists c_1. c \equiv s(n)[c_1]$	$e, s, f \models_F P_F * P'_F \Leftrightarrow \exists f_1, f_2. f \equiv f_1 + f_2 \wedge e, s, f_1 \models_F P_F \wedge e, s, f_2 \models_F P'_F$
$e, s, c \models_C P_C \otimes P'_C \Leftrightarrow \exists c_1, c_2. c \equiv c_1 \otimes c_2$	$e, s, f \models_F \alpha \otimes P_F \Leftrightarrow \exists x, f'. e(\alpha) = x \wedge f \equiv (\nu x)(f') \wedge e, s, f' \models_F P_F$
$\wedge e, s, c_1 \models_C P_C$	$e, s, f \models_F P_F \multimap P'_F \Leftrightarrow \forall f'. e, s, f' \models_F P_F \wedge (f + f') \downarrow \Rightarrow e, s, f + f' \models_F P'_F$
$\wedge e, s, c_2 \models_C P'_C$	$e, s, f \models_F \alpha \neg \otimes P_F \Leftrightarrow \exists x, f'. e(\alpha) = x \wedge f' \equiv (\nu x)(f) \wedge e, s, f' \models_F P_F$
$e, s, c \models_C T \Leftrightarrow \langle c \rangle \equiv \llbracket T \rrbracket s$	$e, s, f \models_F f \Leftrightarrow f \equiv e(f)$
$e, s, c \models_C c \Leftrightarrow c \equiv e(c)$	$e, s, f \models_F B \Leftrightarrow \llbracket B \rrbracket s = \text{true}$
$e, s, c \models_C \langle c \rangle \Leftrightarrow c \simeq e(c)$	$e, s, f \models_F \exists \text{var}. P_F \Leftrightarrow \exists v. e, s[\text{var} \mapsto v], f \models_F P_F$
$e, s, c \models_C B \Leftrightarrow \llbracket B \rrbracket s = \text{true}$	$e, s, f \models_F \exists l \text{var}. P_F \Leftrightarrow \exists v. e[l \text{var} \mapsto v], s, f \models_F P_F$
$e, s, c \models_C @\alpha \Leftrightarrow e(\alpha) \in \text{fn}(c)$	$e, s, f \models_F \forall \alpha. P_F \Leftrightarrow \exists x. x \# e, f \wedge e[\alpha \mapsto x], s, f \models_F P_F$

Fig. 4. Satisfaction Relations of Context Logic for Tree Fragments.

We also use specific formulae for our tree fragment model. The tree context specific formulae are standard and include the variable α which expresses that a tree context is a context hole labeled whose label is the value of variable α . The specific connectives for tree fragments consist of \emptyset_F , describing an empty tree fragment, and $\alpha \leftarrow P_C$, describing a tree context satisfying P_C with hole address given by the value of variable α . This specific formula $\alpha \leftarrow P_C$ is analogous to the atomic formula $n \mapsto n_1, \dots, n_l$ except that we work with hole variables not node identifier variables. We also have existential quantification over location name, tree context and tree fragment variables. Finally, we add the tree context expression formula $@\alpha$ which describes a tree context that contains α free; the analogous formula for tree fragments is derivable.

Definition 10 (Formulae). The formulae of Context Logic for tree fragments include tree context formulae P_C and tree fragment formulae P_F given by:

$P_C ::=$	$P_F ::=$	
$P_C \Rightarrow P_C \mid \text{false}_C$	$P_F \Rightarrow P_F \mid \text{false}_F$	Classical formulae
$\mid \emptyset_C \mid \alpha \mid n[P_C] \mid P_C \otimes P_C$	$\mid P_F * P_F \mid P_F \multimap P_F \mid \alpha \otimes P_F \mid \alpha \neg \otimes P_F$	Structural formulae
$\mid T \mid c \mid \langle c \rangle \mid B \mid @\alpha$	$\mid \emptyset_F \mid \alpha \leftarrow P_C$	Specific formulae
	$\mid f \mid B$	Expression formulae
	$\mid \exists \text{var}. P_F \mid \exists l \text{var}. P_F \mid \forall \alpha. P_F$	Quantification

Notice that the structure of the tree fragment formulae are orthogonal to the structure of the tree context formulae. It is easy to adapt this approach to other data structures such as sequences and terms.

Definition 11 (Satisfaction Relation). Given a logical environment e and a variable store s , the semantics of Context Logic for tree fragments (see Figure 4) is given by two satisfaction relations $e, s, c \models_C P_C$ and $e, s, f \models_F P_F$ defined on tree contexts and tree fragments.

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

$$\begin{aligned}
\text{tree}(P_C) &\triangleq P_C \wedge \neg \exists \alpha. @\alpha & \Diamond P_F &\triangleq \text{true}_F * P_F \\
n &\triangleq n[\emptyset_C] & H\alpha. P_F &\triangleq \forall \alpha. \alpha \otimes P_F \\
o[P_C] &\triangleq \exists m. m[P_C]
\end{aligned}$$

Formula $\text{tree}(P_C)$ describes a complete tree. Formula n allows us to drop the need to mention when a subtree is empty. Formula $\circ[P_C]$ allows us to drop the identifier of a node. Formula $\diamond P_F$ allows us to express that, somewhere in the tree fragment, P_F holds. Finally, the hiding quantification, $\text{H}\alpha$, is shorthand for revelation over a fresh label.

Example 6 (Context Logic Examples).

- (a) The tree fragment formula $\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\delta]$ describes a tree fragment consisting of a node n with address α and context hole γ and node m with address β and context hole δ ; the n and m are non-equal. Now consider the fragment formula:

$$\omega \leftarrow n[\gamma] \otimes m[\delta] \Leftrightarrow \alpha, \beta \textcircled{\text{R}} (\omega \leftarrow \alpha \otimes \beta * \alpha \leftarrow n[\gamma] * \beta \leftarrow m[\delta])$$

The first formula states that the nodes n and m are siblings with address ω . The second formula states that the node n has address α , the node m has address β and the holes α and β are siblings with address ω . The labels α and β are revealed in the fragment and thus these two formulae are equivalent. The separation connective $*$ allows us to add more pieces to the tree fragment and the revelation connective $\textcircled{\text{R}}$ allows us to link up, and break apart, these pieces.

- (b) The tree fragment formula $\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\text{tree}(c)]$ describes a tree fragment consisting of a single node n at address α and a complete tree with top node m at address β . This formula is the precondition of the small axiom of `appendChild`. In particular, due to the satisfaction relation for $*$, we know that node n cannot appear in the tree with top node m without violating the unique location name requirement of the tree fragment. This style of formula allows us to capture the tree fragments required for the successful running of the move subtree commands of our update language, elegantly expressing both the case where the trees at n and m are disjoint and the case where n is an ancestor of m .
- (c) The tree fragment formula $\exists c. \text{H}\alpha. ((\alpha \leftarrow n[\emptyset_C] \rightarrow (\alpha \rightarrow \textcircled{\text{R}} P_F)) * \alpha \leftarrow n[\text{tree}(c)])$ describes a tree fragment which can be split into a complete tree with top node n at address α and a tree fragment, that when extended by node n with the empty tree beneath it satisfies some property P_F . The use of $\textcircled{\text{R}}$ hides the label α allowing us to pull the tree with top node n out of the larger fragment. The use of $\rightarrow \textcircled{\text{R}}$ describes putting the extracted tree fragment back into the larger fragment. If this tree fragment has had the subtree at n removed then the property P_F must now hold. This formula is the weakest precondition of the `deleteSubtree(n)` command.

5 Local Hoare Reasoning

We use the logic defined in Section 4 to provide local Hoare reasoning about programs written in the language defined in Definition 6. First we give a fault avoiding partial correctness interpretation of local Hoare triples following [18].

$\{\emptyset_F\}$	skip	$\{\emptyset_F\}$
$\{\emptyset_F \wedge (n = n_0)\}$	$n := N$	$\{\emptyset_F \wedge (n = N[n_0/n])\}$
$\{\emptyset_F \wedge (t = t_0)\}$	$t := T$	$\{\emptyset_F \wedge (t = T[t_0/t])\}$
$\{\alpha \leftarrow n[\text{tree}(c)]\}$	$t := \text{copy}(n)$	$\{\alpha \leftarrow n[\text{tree}(c)] \wedge (t = \langle n[\text{tree}(c)] \rangle)\}$
$\{\alpha \leftarrow m[\beta \otimes n[\delta] \otimes \gamma] \wedge (n' = n_0)\}$	$n' := \text{getUp}(n)$	$\{\alpha \leftarrow m[\beta \otimes n_{[n_0/n']}[\delta] \otimes \gamma] \wedge (n' = m)\}$
$\{\alpha \leftarrow n[\delta] \otimes m[\beta] \wedge (n' = n_0)\}$	$n' := \text{getRight}(n)$	$\{\alpha \leftarrow n_{[n_0/n']}[\delta] \otimes m[\beta] \wedge (n' = m)\}$
$\{\alpha \leftarrow m[\beta \otimes n[\delta]] \wedge (n' = n_0)\}$	$n' := \text{getRight}(n)$	$\{\alpha \leftarrow m[\beta \otimes n_{[n_0/n']}[\delta]] \wedge (n' = \text{null})\}$
$\{\alpha \leftarrow n[\delta \otimes m[\beta]] \wedge (n' = n_0)\}$	$n' := \text{getLast}(n)$	$\{\alpha \leftarrow n_{[n_0/n']}[\delta \otimes m[\beta]] \wedge (n' = m)\}$
$\{\alpha \leftarrow n[\emptyset_C] \wedge (n' = n_0)\}$	$n' := \text{getLast}(n)$	$\{\alpha \leftarrow n_{[n_0/n']}[\emptyset_C] \wedge (n' = \text{null})\}$
$\{\alpha \leftarrow n[\beta]\}$	deleteNode(n)	$\{\alpha \leftarrow \beta\}$
$\{\alpha \leftarrow n[\text{tree}(c)]\}$	deleteSubtree(n)	$\{\alpha \leftarrow n[\emptyset_C]\}$
$\{\alpha \leftarrow n[\beta]\}$	insertNodeAbove(n)	$\{\alpha \leftarrow \circ[n[\beta]]\}$
$\{\alpha \leftarrow n[\beta]\}$	insertRight(n, T)	$\{\alpha \leftarrow n[\beta] \otimes T\}$
$\{\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\delta]\}$	moveNodeAbove(n, m)	$\{\alpha \leftarrow m[n[\gamma]] * \beta \leftarrow \delta\}$
$\{\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\delta]\}$	appendNode(n, m)	$\{\alpha \leftarrow n[\gamma \otimes m[\emptyset_C]] * \beta \leftarrow \delta\}$
$\{\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\text{tree}(c)]\}$	appendSub(n, m)	$\{\alpha \leftarrow n[\gamma \otimes \text{tree}(c)] * \beta \leftarrow m[\emptyset_C]\}$

Fig. 5. Small Axioms for the Tree Update Language.

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

$$\{P_F\} \mathbb{C} \{Q_F\} \quad \Leftrightarrow \quad \forall e, s, f. \text{free}(\mathbb{C}) \cup \text{free}(P) \cup \text{free}(Q) \subseteq \text{dom}(s) \wedge e, s, f \models_F P_F \\ \Rightarrow \mathbb{C}, s, f \rightsquigarrow \text{fault} \wedge \forall s', f'. \mathbb{C}, s, f \rightsquigarrow s', f' \Rightarrow e, s', f' \models_F Q_F$$

Definition 14 (Small Axioms). The Small Axioms are given in Figure 5.

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

$$\begin{array}{c} \text{F} \quad \text{L} \quad \text{E} \quad : \quad \text{A} \quad \text{V} \quad \text{E} \quad : \\ \frac{\{P_F\} \mathbb{C} \{Q_F\}}{\{\mathcal{W}\alpha. P_F\} \mathbb{C} \{\mathcal{W}\alpha. Q_F\}} \quad \frac{\{P_F\} \mathbb{C} \{Q_F\}}{\{\exists n. P_F\} \mathbb{C} \{\exists n. Q_F\}} \quad n \notin \text{Free}(\mathbb{C}) \\ \\ \text{R} \quad \quad \text{F} \quad : \quad \quad \text{S} \quad \quad \text{F} \quad : \\ \frac{\{P_F\} \mathbb{C} \{Q_F\}}{\{\alpha \textcircled{R} P_F\} \mathbb{C} \{\alpha \textcircled{R} Q_F\}} \quad \frac{\{P_F\} \mathbb{C} \{Q_F\}}{\{P_F * R_F\} \mathbb{C} \{Q_F * R_F\}} \quad \text{Mod}(\mathbb{C}) \cap \text{Free}(R_F) = \{\} \end{array}$$

The Auxiliary Variable Elimination and Separation Frame rules are standard from Separation Logic. The Revelation Frame rule is the natural consequence of having restriction in the model. The Fresh Label Elimination rule is analogous to the Auxiliary Variable Elimination rule.

Our reasoning system is sound. The weakest preconditions of our tree update commands (given in Figure 6) are derivable; proof in full paper [8]. This means that our local Hoare reasoning is complete for straight line code.

6 Examples

We provide specifications for each of the example programs given in section 2. We make the assumption that all locally defined variables, such as `temp` in the `appendChild` program, are disjoint from all other program variables.

	$\{P_F\}$	skip	$\{P_F\}$
	$\{\exists n_0. (n = n_0) \wedge P_F[N/n]\}$	$n := N$	$\{P_F\}$
	$\{\exists t_0. (t = t_0) \wedge P_F[T/t]\}$	$t := T$	$\{P_F\}$
	$\{\exists c. \text{Ha}. \phi \alpha \leftarrow n[\text{tree}(c)] \wedge (\alpha \rightarrow P_F[n[\text{tree}(c)]]/t)\}$	$t := \text{copy}(n)$	$\{P_F\}$
$\{\exists m, n_0. \text{Ha}. \beta, \gamma, \delta. \phi \alpha \leftarrow m[\beta \otimes n[\delta]] \otimes \gamma \wedge (n' = n_0) \wedge (\alpha, \beta, \gamma, \delta \rightarrow P_F[m/n'])\}$		$n' := \text{getUp}(n)$	$\{P_F\}$
$\{\exists m, n_0. \text{Ha}. \beta, \gamma, \delta. \phi \alpha \leftarrow m[\beta \otimes n[\delta]] \wedge (n' = n_0) \wedge (\alpha, \beta, \gamma, \delta \rightarrow P_F[m/n'])\}$		$n' := \text{getRight}(n)$	$\{P_F\}$
$\{\exists m, n_0. \text{Ha}. \beta, \gamma, \delta. \phi \alpha \leftarrow m[\delta \otimes m[\beta]] \wedge (n' = n_0) \wedge (\alpha, \beta, \gamma, \delta \rightarrow P_F[m/n'])\}$		$n' := \text{getLast}(n)$	$\{P_F\}$
$\{\exists m, n_0. \text{Ha}. \beta, \gamma, \delta. \phi \alpha \leftarrow m[\delta \otimes m[\beta]] \wedge (n' = n_0) \wedge (\alpha, \beta, \gamma, \delta \rightarrow P_F[m/n'])\}$		$\text{deleteNode}(n)$	$\{P_F\}$
$\{\exists c. \text{Ha}. ((\alpha \leftarrow n[\mathcal{O}_C] \rightarrow (\alpha \rightarrow P_F)) * \alpha \leftarrow n[\text{tree}(c)])\}$		$\text{deleteSubtree}(n)$	$\{P_F\}$
$\{\text{Ha}. \beta. ((\alpha \leftarrow n[\beta] \rightarrow (\alpha, \beta \rightarrow P_F)) * \alpha \leftarrow n[\beta])\}$		$\text{insertNodeAbove}(n)$	$\{P_F\}$
$\{\text{Ha}. \beta. ((\alpha \leftarrow n[\beta] \otimes T \rightarrow (\alpha, \beta \rightarrow P_F)) * \alpha \leftarrow n[\beta])\}$		$\text{insertRight}(n, T)$	$\{P_F\}$
$\{\text{Ha}. \beta, \gamma, \delta. (((\alpha \leftarrow m[n[\gamma]] * \beta \rightarrow \delta) \rightarrow (\alpha, \beta, \gamma, \delta \rightarrow P_F)) * (\alpha \leftarrow n[\gamma] * \beta \rightarrow m[\delta]))\}$		$\text{moveNodeAbove}(n, m)$	$\{P_F\}$
$\{\text{Ha}. \beta, \gamma, \delta. (((\alpha \leftarrow n[\gamma \otimes m[\mathcal{O}_C]] * \beta \rightarrow \delta) \rightarrow (\alpha, \beta, \gamma, \delta \rightarrow P_F)) * (\alpha \leftarrow n[\gamma] * \beta \rightarrow m[\delta]))\}$		$\text{appendNode}(n, m)$	$\{P_F\}$
$\{\exists c. \text{Ha}. \beta, \gamma. (((\alpha \leftarrow n[\gamma \otimes \text{tree}(c)] * \beta \rightarrow m[\mathcal{O}_C]) \rightarrow (\alpha, \beta, \gamma \rightarrow P_F)) * (\alpha \leftarrow n[\gamma] * \beta \rightarrow m[\text{tree}(c)]))\}$		$\text{appendSub}(n, m)$	$\{P_F\}$

Fig. 6. Weakest Preconditions of the atomic commands given in Figure 3.

appendChild: In example 1 of section 2 we gave the program `appendChild`. Its specification and derivation are:

$$\begin{array}{l}
\{\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\text{tree}(c)]\} \\
\{\text{H}\delta. \alpha \leftarrow n[\gamma] * \beta \leftarrow m[\delta] * \delta \leftarrow \text{tree}(c)\} \\
\text{insertRight}(m, \mathcal{O}[\mathcal{O}_T]); \\
\{\text{H}\delta. \alpha \leftarrow n[\gamma] * \beta \leftarrow m[\delta] \otimes \mathcal{O}[\mathcal{O}_T] * \delta \leftarrow \text{tree}(c)\} \\
\text{temp} := \text{getLeft}(m); \\
\{\text{H}\delta. \alpha \leftarrow n[\gamma] * \beta \leftarrow m[\delta] \otimes \text{temp}[\mathcal{O}_C] * \delta \leftarrow \text{tree}(c)\} \\
\{\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\text{tree}(c)] \otimes \text{temp}[\mathcal{O}_C]\} \\
\{\text{H}\delta, \epsilon. \alpha \leftarrow n[\gamma] * \beta \leftarrow \delta \otimes \epsilon * \delta \leftarrow m[\text{tree}(c)] * \epsilon \leftarrow \text{temp}[\mathcal{O}_C]\} \\
\text{appendSub}(\text{temp}, m); \\
\{\text{H}\delta, \epsilon. \alpha \leftarrow n[\gamma] * \beta \leftarrow \delta \otimes \epsilon * \delta \leftarrow m[\mathcal{O}_C] * \epsilon \leftarrow \text{temp}[\text{tree}(c)]\} \\
\text{appendNode}(n, m); \\
\{\text{H}\delta, \epsilon. \alpha \leftarrow n[\gamma \otimes m[\mathcal{O}_C]] * \beta \leftarrow \delta \otimes \epsilon * \delta \leftarrow \mathcal{O}_C * \epsilon \leftarrow \text{temp}[\text{tree}(c)]\} \\
\text{appendSub}(m, \text{temp}); \\
\{\text{H}\delta, \epsilon. \alpha \leftarrow n[\gamma \otimes m[\text{tree}(c)]] * \beta \leftarrow \delta \otimes \epsilon * \delta \leftarrow \mathcal{O}_C * \epsilon \leftarrow \text{temp}[\mathcal{O}_C]\} \\
\text{deleteNode}(\text{temp}) \\
\{\text{H}\delta, \epsilon. \alpha \leftarrow n[\gamma \otimes m[\text{tree}(c)]] * \beta \leftarrow \delta \otimes \epsilon * \delta \leftarrow \mathcal{O}_C * \epsilon \leftarrow \mathcal{O}_C\} \\
\{\alpha \leftarrow n[\gamma \otimes m[\text{tree}(c)]] * \beta \leftarrow \mathcal{O}_C\}
\end{array}$$

The `appC(n, m)` program below has equivalent behavior to `appendChild(n, m)` modulo renaming of the tree at m , and an analogous specification:

$$\begin{array}{l}
\text{appC}(n, m) \triangleq t := \text{copy}(m); \\
\text{deleteTree}(m); \\
\text{insertLast}(n, t)
\end{array}
\quad
\begin{array}{l}
\{\alpha \leftarrow n[\gamma] * \beta \leftarrow m[\text{tree}(c)]\} \\
\text{appC}(n, m) \\
\{\alpha \leftarrow n[\gamma \otimes m[\text{tree}(c)]] * \beta \leftarrow \mathcal{O}_C\}
\end{array}$$

The derivation of this specification is similar to the derivation of `appendChild` given above. In multi-holed Context logic we could give small specifications for each of the atomic commands used to construct the `appC` program. However, we cannot provide a small specification for `appC` directly from the small axioms of these atomic commands. As we discussed in the introduction, we instead must use a specification with a context variable C to describe the linking context between nodes n and m .

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

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

The derivations of these specifications are shown in Figure 7.

simple derivation:

```

{α ← n[m[β] ⊗ γ]}
insertNodeAbove(n);
{α ← o [n[m[β] ⊗ γ]]}
temp := getUp(n);
{α ← temp[n[m[β] ⊗ γ]]}
first := getFirst(n);
{α ← temp[n[first[β] ⊗ γ]] ∧ (m = first)}
moveNodeAbove(first, n);
{α ← temp[n[first[β] ⊗ γ] ∧ (m = first)}
moveNodeAbove(temp, first);
{α ← first[temp[n[β] ⊗ γ]] ∧ (m = first)}
deleteNode(temp)
{α ← first[n[β] ⊗ γ] ∧ (m = first)}
{α ← m[n[β] ⊗ γ]}

```

nodeSwap derivation:

```

{α ← n[γ] * β ← m[δ]}
insertNodeAbove(n);
{α ← o [n[γ]] * β ← m[δ]}
temp := getUp(n);
{α ← temp[n[γ]] * β ← m[δ]}
moveNodeAbove(m, n);
{α ← temp[γ] * β ← n[m[δ]]}
moveNodeAbove(temp, m);
{α ← m[temp[γ]] * β ← n[δ]}
deleteNode(temp)
{α ← m[γ] * β ← n[δ]}

```

nodeCycle derivation:

```

{α ← n[m[β] ⊗ γ] l[δ]}
first := getFirst(n);
{ α ← n[first[β] ⊗ γ] l[δ] }
  ∧ (first = m)
last := getLast(n);
{ α ← n[first[β] ⊗ γ] last[δ] }
  ∧ (first = m) ∧ (last = l)
nodeSwap(n, last);
{ α ← last[first[β] ⊗ γ] n[δ] }
  ∧ (first = m) ∧ (last = l)
nodeSwap(n, first)
{ α ← last[n[β] ⊗ γ] first[δ] }
  ∧ (first = m) ∧ (last = l)
{α ← l[n[β] ⊗ γ] m[δ]}

```

Fig. 7. Derivations of the specifications for simple, nodeSwap and nodeCycle.

Hierarchical Queue: The specification and derivation of the queuePop program from example 5 of section 2 are:

```

{α ← n[m[tree(c)] ⊗ γ ⊗ i[β]]}
next := getFirst(n);
{α ← n[next[tree(c)] ⊗ γ ⊗ i[β]] ∧ (next = m)}
info := getLast(n);
{α ← n[next[tree(c)] ⊗ γ ⊗ info[β]] ∧ (next = m) ∧ (info = i)}
nodeSwap(n, info);
{α ← info[next[tree(c)] ⊗ γ ⊗ n[β]] ∧ (next = m) ∧ (info = i)}
nodeSwap(next, info);
{α ← next[info[tree(c)] ⊗ γ ⊗ n[β]] ∧ (next = m) ∧ (info = i)}
appendChild(next, info)
{α ← next[γ ⊗ n[β] ⊗ info[tree(c)]] ∧ (next = m) ∧ (info = i)}
{α ← m[γ ⊗ n[β] ⊗ i[tree(c)]]}

```

7 Conclusion

We have shown how to give small axioms for commands such as the `appendChild` command, by developing Context Logic reasoning for tree fragments. It is straightforward to transfer the techniques developed here to Featherweight DOM [7]. For this paper, we have worked with the intuitive understanding of what it means for command axioms to be small. With Raza, Gardner has developed the formal definitions of footprints and small specifications for abstract local functions using Abstract Separation Logic [14]. It would be interesting to extend this abstract theory to the tree fragments and reasoning studied here, and prove that the axioms really are small.

We believe the results presented here form a pivotal step in the development of Context Logic reasoning. The key point about Context Logic is that it reasons about structured data at the same level of abstraction as the data itself. Our previous work used various forms of separating application, which were appropriate for the applications we had in mind. Here, we move nearer to Separation Logic reasoning. We use the separating conjunction for reasoning about disjoint tree fragments, and the revelation connectives and freshness quantification for reasoning about restriction. This reasoning style means that we can pull out different tree fragments from the working tree, update them, and put them back again in any undetermined order. The ideas in this paper provides us with the technology to extend our reasoning to concurrent tree update following O'Hearn's work on concurrent Separation Logic [1],[13],[16].

References

1. S. Brookes. A semantics for concurrent separation logic. *Theor. Comput. Sci.*, 375, 2007.
2. C. Calcagno, T. Dinsdale-Young, and P. Gardner. Adjoint elimination in context logic for trees. In *APLAS*, 2007.
3. C. Calcagno, P. Gardner, and U. Zarfaty. Context logic and tree update. In *POPL*, 2005.
4. L. Cardelli and A. D. Gordon. Ambient logic. 2006.
5. M. J. Gabbay and A. M. Pitts. A new approach to abstract syntax with variable binding.
6. P. Gardner, G. Smith, M. Wheelhouse, and U. Zarfaty. Dom: Towards a formal specification. In *Plan-X: Programming Language Techniques for XML*, 2008.
7. P. Gardner, G. Smith, M. Wheelhouse, and U. Zarfaty. Local hoare reasoning about dom. In *PODS: Symposium on Principles of Database Systems*, 2008.
8. P. Gardner and M. Wheelhouse. Small specifications for tree update, 2009. <http://www.doc.ic.ac.uk/~mjw03/PersonalWebpage/papers.html>.
9. P. Gardner and U. Zarfaty. Integrated reasoning about high-level tree update and a low-level implementation. submitted to publication.
10. S. Ishtiaq and P. W. O’Hearn. Bi as an assertion language for mutable data structures. In *POPL*, pages 14–26, 2001.
11. R. Milner. A calculus of mobile processes, parts. *I and II. Information and Computation*, 100:1–77, 1992.
12. P. O’Hearn, J. Reynolds, and H. Yang. *Local Reasoning about Programs that Alter Data Structures*, volume 2142. January 2001.
13. P. W. Ohearn. Resources, concurrency and local reasoning. In *Theoretical Computer Science*, pages 49–67. Springer, 2004.
14. M. Raza and P. Gardner. Footprints in local reasoning. In *FoSSaCS ’08: Proceedings of the 11th International Conference on Foundations of Software Science and Computation Structures*, volume 4962, pages 201–215, London, UK, 2008. Springer.
15. G. Smith. Providing a formal specification for dom core level 1, 2009. PhD Thesis. Ongoing work.
16. V. Vafeiadis and M. Parkinson. A marriage of rely/guarantee and separation logic. In *In 18th CONCUR*, pages 256–271. Springer, 2007.
17. W3C. Dom: Document object model. W3C recommendation, 2005. <http://www.w3.org/DOM/>.
18. H. Yang and P. W. O’Hearn. A semantic basis for local reasoning. In *FoSSaCS ’02: Proceedings of the 5th International Conference on Foundations of Software Science and Computation Structures*, pages 402–416, London, UK, 2002. Springer-Verlag.